

MAAWG IPv6 Training for Senders and Others

Joe St Sauver, Ph.D. (joe@oregon.uoregon.edu)

MAAWG, Washington DC, October 4th-6th, 2010

www.uoregon.edu/~joe/maawg-senders-ipv6-training/

Disclaimer: All opinions expressed in this talk are strictly my own, and do not necessarily represent the opinions of any other entity. This talk is provided in a detailed written form to insure accessibility, and for ease of web indexing.

Welcome to the DC Area and MAAWG's IPv6 Training for Senders (and Others)!

- We're excited to be with you here today to talk for about seventy five minutes about what you might want to know about IPv6 if you're a sender.
- Just to get started, let's go around the room and have each person briefly say:
 - who they are
 - what they do
 - the company you're with
 - a little about your interest in/experience with IPv6
 - your technical level (non-technical, semi-technical, hard-core geek, or whatever)

PART I.
Basic IPv6 For Senders

1. Do Mail Senders *Really* Need IPv6 Today?

If You're A Mail Sender, Do You Really Need IPv6 Today? Well, No.

- You obviously already have IPv4 address space.
- The sites you're sending to, with perhaps just a few exceptions, will not have IPv6-enabled mail transfer agents (MTAs), so even if you wanted to talk to them via IPv6, they may not be ready for you to do so.
- If the sites you're sending to do have IPv6-enabled MTAs, in virtually every case those MTAs will be actually be "dual stack," e.g., they'll support both IPv6 and IPv4 transport, they won't be IPv6-only MTAs. So, if you can already reach a site via IPv4, why bother trying IPv6 instead, particularly if some of that IPv6 connectivity may be tunnelled and indirect, slow, lossy, or otherwise lower quality than also-available IPv4 links?

But, Even If You Don't Need To Move to IPv6 As A Sender, Others Do

- For example, within a year or two, there will not be any more IPv4 address space to allocate to growing ISPs working hard to hook up new customers, or ISPs trying to accommodate a growing number of devices/customer.
- Alternatives to IPv6, such as using Network Address Translation (NAT) with private address space, pose some really ugly operational challenges for ISPs (particularly for things like tracking down malware-infested customers living behind a NAT box)
- Software and hardware vendors that need to service a market that's moving to IPv6 also need to "make the IPv6 leap" and enable IPv6 in their products as a result of their customers' emerging requirements.

Bottom Line

- Since you're already here, you might as well at least learn a little about IPv6, and maybe even give it a try when you get back home
- The process doesn't have to be painful, and if you don't like it, as a sender you can always "back it back out" and go back to just using IPv4.
- The first part of this talk will be for those of you who are busy and goal oriented: we'll start by looking at what you'd need to do if you did want to begin sending email traffic from an IPv6-enabled server on an IPv6-enabled network.

2. IPv6 and SMTP

Email Is The "Forgotten" Application of IPv6

- While many people are very excited about the thought of using IPv6 for the **web**, for some reason there seems to be a lot less excitement about using IPv6 for **email**.
- Thus, while many mainstream mail software products support IPv6, relatively few mail administrators apparently bother to enable IPv6 support.
- But some sites **ARE** deploying IPv6-accessible mail servers right now. For example...

Sample Institutional IPv6 Enabled MX

```
% dig ucla.edu mx +short  
5 smtp.ucla.edu.
```


```
% dig smtp.ucla.edu a +short  
169.232.46.240  
169.232.46.241  
169.232.46.242  
169.232.46.244  
etc.
```

```
% dig smtp.ucla.edu aaaa +short  
2607:f010:3fe:302:1013:72ff:fe5b:60c3  
2607:f010:3fe:102:101c:23ff:febe:116e  
2607:f010:3fe:102:101c:23ff:febf:cfa7  
2607:f010:3fe:102:101c:23ff:fed0:918c  
etc.
```

Examples of Other IPv6 Enabled Domains/MXs

brocade.com (mx10.brocade.com)
maine.edu (mail-relay.maine.edu)
vt.edu (inbound.smtp.vt.edu)
iij.ad.jp (omgi.iij.ad.jp)
ams-ix.net (betonmix.noc.ams-ix.net)
apnic.net (hoisin.apnic.net)
arin.net (smtp2.arin.net)
dren.net (drenmail.dren.net)
es.net (mail.es.net)
he.net (he.net)
jpnict.net (mx15.nic.ad.jp)
lacnic.net (mail.lacnic.net)
ripe.net (postgirl.ripe.net)
icann.org (pechora2.icann.org)
ietf.org (mail.ietf.org)
isc.org (mx.ams1.isc.org)
kth.se (mx.kth.se)

Code To Check for IPv6 MX Records



Email and IPv6

But what about email? How many servers can receive email via IPv6? We took a sample of 500,000 domains. Unfortunately we cannot use the list from [Alexa](#) as they are web sites and not domains with mail servers. Instead, we sampled 500,000 email domains known to [Genius.com](#) and went to look for their [MX records](#). For each MX record we checked if the host has an [AAAA record](#) (reachable via IPv6).

We used the simple program below against a csv file of domains:

```
<?php
$file=$argv[1];
$f=fopen($file,"r");
$buffer = fgets($f, 4096);
$i=1;
while (!feof($f)) {
    $buffer = fgets($f, 4096);
    $domain=substr($buffer,1,-2);
    echo $i."|".$domain."|";

    //look for MX record
    $mxhosts=array();
    $foundMX=getmxrr($domain,&$mxhosts);
    if ($foundMX) {
        //loop through MX records
        $ipv6=false;
        foreach($mxhosts as $host) {
            if (checkdnsrr($host,"AAAA")) {
                $ipv6=true;
                break;
            }
        }
    }
}
```

<-- must be php5 (for getmxrr)

<-- (\$buffer,0,-1)

<-- (\$domain,&\$mxhosts)

[continues; see the original site for the full code] 12

Important Note:

- While we may be curious to find out who's currently running with SMTP enabled over IPv6, we do NOT need to manually keep track of those sites, NOR do we need to change any of the addresses we mail to once both the sender and receiver are IPv6 enabled and IPv6 accessible.
- In most cases, mail servers will simply **automatically select IPv6 transport** if both the sender and the receiver support IPv6.
- As a sender, then, you simply need to:
 - get IPv6 connectivity (and IPv6 address space)
 - enable IPv6 for your server's operating system
 - enable IPv6 for your MTA
- Go over ALL the things you'll need to do before you start ordering connectivity and start modifying gear, etc.!

3. Obtaining Native IPv6 Connectivity And Address Space

Add IPv6 Internet Transit Connectivity

- This task is largely a business office/financial one.
- Contact your account manager at your current colo or network service provider and tell him/her that you'd like to add IPv6 transit to your existing IPv4 connectivity.
- While you're at it, ask them for IPv6 address space, too. You'll probably get a /48 worth of IPv6 address space by default (we'll talk about that later); a /48 should be fine.
- If your account manager doesn't know what IPv6 transit connectivity is, press them to ask their sales engineering support person (or otherwise escalate your question).
- If they check and they really don't offer native IPv6 (or they do, but not at your location/not for you, etc.), see the list of providers who do have IPv6 transit available at <http://www.sixxs.net/faq/connectivity/?faq=ipv6transit>

How Much IPv6 Connectivity Do I Need?

- Your initial IPv6 connectivity requirements will probably be relatively modest. Ideally, your IPv6 connectivity and your IPv4 connectivity will share your existing bandwidth and links, and if you're able to do that, then you won't end up needing to purchase separate IPv6 connectivity.
- If you do need to purchase separate IPv6 connectivity, perhaps because your current provider isn't able to deliver IPv6 connectivity to you, I'd suggest starting modestly, maybe with IPv6 transit capacity at no more than 10-15% of your current IPv4 transit bandwidth level.
- You'll probably have more capacity than you'll initially need, but you'll find that usage will grow over time. Be sure that whatever plan you select gives you flexibility to adjust your IPv6 transit capacity if you need to do so.

“What About Just Using a Free Tunnel?”

- An alternative to getting IPv6 transit connectivity from a network service provider is getting a free IPv6 tunnel from a tunnel broker, such as Hurricane Electric (see <http://tunnelbroker.net/>).
- Their free tunnel service, and similar tunnel services, are great as far as they go, but IPv6 tunnel services tend to be oriented toward developers and experimenters, and since you're a sender, you're going to want production grade native IPv6 service instead.
- If you can't get native IPv6 connectivity, I'd probably wait until you can do so before “taking the IPv6 leap.”

Networking Equipment: Routers and Switches

- Depending on your setup, you may have your own customer premises equipment (CPE), such as a Cisco or Juniper border router, or you may just get a 100Mbps or gigabit ethernet handed to you from your provider, perhaps fanned out across some layer 2 ethernet switches.
- Virtually all layer 2 ethernet switches will transparently pass IPv6 traffic without requiring any tweaking or adjustment.
- If you have your own layer 3 device, such as a Cisco or Juniper router, you will need to have your network engineer enable IPv6 on it.
- Note that some older-generation Ciscos may not be able to process IPv6 at full wirespeed; you should be planning to upgrade or replace geriatric equipment of that sort.₁₈

Networking Equipment: Firewalls

- Some sites may have an IPv4 hardware or software firewall configured in front of their servers. Two things to be aware of...
- First, some firewalls may not be “IPv6 aware” and may just drop any IPv6 packets they see. Those firewalls may need to be upgraded, replaced or bypassed.
- Second, assuming your existing firewall is able to handle IPv6, and you want similar protection for connections over IPv6 as for IPv4, don't forget to create rules allowing desired IPv6 traffic and blocking remaining IPv6 traffic! Your IPv4 rules will NOT automatically be extended to encompass IPv6 traffic, and so, by default, you may either unexpectedly block legitimate IPv6 traffic, or unexpectedly permit unwanted IPv6 traffic.

Networking Equipment: DNS Servers

- Most senders will routinely use two sorts of DNS servers: recursive resolvers and authoritative name servers.
 - Authoritative name servers answer DNS queries for your domains
 - Recursive resolvers resolve all other domain names
- Both sort of name servers will need to support IPv6 resource records (such as AAAA records), however neither needs to support access to the name servers over IPv6 (for now, you can just continue to access your recursive resolvers and authoritative name servers over IPv4).
- If your local recursive resolvers don't support IPv6 and you don't want to upgrade them to do so, one alternative would be to use Google's intentionally open recursive resolvers at 8.8.8.8 and 8.8.4.4 -- they do support IPv₂6.

Networking Equipment: Other Stuff

- Unlike ISPs with end users, senders don't need to worry about things like dynamically assigning IPv6 addresses to end users/customers; we assume that you'll be manually/statically assigning all IPv6 server addresses instead. Therefore, you don't need to worry about DHCPv6 or stateless autoconfiguration or any of the associated IP address assignment gyrations.
- Similarly, if you do central syslogging, or centralized device monitoring and management, or centralized NTP (time service) currently, we assume that you will continue to do those things over IPv4 -- for now, the only things that will be talking over IPv6 transport are your MTAs.

Your New IPv6 Address Space

- If you're a sender who currently maybe has an IPv4 /24 with 255 addresses, and you suddenly get a /48 worth of IPv6 space (e.g., $2^{(128-48)}$, or 1,208,925,819,614,629,174,706,176 IPv6 addresses), you may be overwhelmed by all your new address space.
- It may help if you think of that as "just" 65,536 subnets, with each subnet being 64 bits long (or having 2^{64} or 18,446,744,073,709,551,616 addresses). You will likely only need one of those, but it's good to have room to grow. :-)
- For now, let's just assume that your network engineer (or your network service provider/colo provider) will recommend appropriate IPv6 host IP addresses from your IPv6 range, and we'll also assume that he or she will tell you your default upstream IPv6 gateway address.

4. Enabling IPv6 on Your Servers

Most Modern Operating Systems Ship “IPv6 Ready”

- In general, most modern operating systems are IPv6 ready; if you’re configuring a mail server, you just need to enable IPv6 and supply basic configuration information (such as the server’s static IPv6 address).
- Q. “Help! My operating system doesn’t have built in support for IPv6! Should I patch it or something?”
A. If your operating system requires a patch to support IPv6, it is ancient and riddled with unpatched (and unpatchable) security issues. Before spending time experimenting with IPv6, first things first: upgrade your operating system to the current production release of your O/S (it **will** have native support for IPv6).

Enabling IPv6 with A Static IPv6 Address

- Recipes to enable IPv6 and assign a static IPv6 addr for common OS:

-- Redhat/CentOS:

www.cyberciti.biz/faq/rhel-redhat-fedora-centos-ipv6-network-configuration/

-- SuSE Linux:

www.cyberciti.biz/faq/configuring-ipv6-in-sles10-opensuse-linux/

-- Ubuntu Linux:

www.cyberciti.biz/faq/ubuntu-ipv6-networking-configuration/

-- FreeBSD and friends:

www.cyberciti.biz/faq/freebsd-configure-ipv6-networking-static-ip-address/

-- Windows Server 2008/R2

technet.microsoft.com/en-us/library/cc732106.aspx

Mac user? Just set a static IP in System Preferences

-> Network -> Configure -> Configure IPv6 -> Manually

Example: Enabling IPv6 in Redhat with A Static IP

- In `/etc/sysconfig/network`

```
NETWORKING_IPV6=yes
```

- In `/etc/sysconfig/network-scripts/ifcfg-eth0`

```
IPV6INIT=yes
```

```
IPV6ADDR=your_servers_IPv6_address_here
```

```
IPV6_DEFAULTGW=your_servers_default_gw_here
```

- `# service network restart`

5. Enabling IPv6 in Your MTA

IPv6 Support In Mainstream Email Software Products

- Virtually all modern mail transfer agents support IPv6; a lack of IPv6 MTA software support is typically NOT an issue...
 - Exchange 2007 SP1 (only under Windows Server 2008, and only with both IPv4 and IPv6 enabled); see <http://technet.microsoft.com/en-us/library/bb629624.aspx>
 - Exim (http://www.exim.org/exim-html-current/doc/html/spec_html/ch04.html at section 9)
 - Postfix (http://www.postfix.org/IPV6_README.html)
 - Qmail (via Qsmtp, see <http://opensource.ssf-tec.de/Qsmtp/>)
 - Sendmail (see the Sendmail Installation and Operation Guide, <http://www.sendmail.org/doc/sendmail-current/doc/op/op.pdf>)
- What about Procmail as a local mail delivery agent? Umm, well, unfortunately see <http://www.procmail.org/todo.html>

What It Takes to Enable IPv6 in postfix

- For now, we'll assume that you have a ready-to-go IPv6-enabled network, and the only thing holding you up is a non-IPv6 aware MTA
- Let's assume you want to use Postfix. Get postfix 2.7 (or whatever's the latest production code) from www.postfix.org/download.html
- Review http://www.postfix.org/IPV6_README.html
- When configuring for IPv6, in `/etc/postfix/main.cf`, set `inet_protocols = ipv6, ipv4` (if you're dual stacking)
- Also include in `/etc/postfix/main.cf` the address you want to use for outgoing IPv6 SMTP connections:
`smtp_bind_address6 = 2001:468:d01:d6::80df:d617` <-- sample only!
- Check your config and start postfix; typically:
`/usr/sbin/postfix check`
`/usr/sbin/postfix start`
- Confirm that you can connect OK to your IPv6 smtpd:
`% telnet 2001:468:d01:d6::80df:d617 25` <-- sample only!
`quit`

IPv6 and DNS Blocklists

- DNS blocklists, such as those offered by Spamhaus, are a key anti-abuse tool in today's IPv4-dominated Internet, directly blocking spam while also “encouraging” ISPs to employ sound anti-abuse practices.
- Virtually all sites that use DNS-based blocklists rely on rblDNS (see www.corpit.ru/mjt/rblDNS/rblDNS.8.html). rblDNS does NOT support IPv6 records at this time. :-(
- Spamhaus does not maintain any substantive IPv6 blocklists; Spamhaus has, however, just recently announced a new IPv4 and IPv6 whitelist (see <http://www.spamhauswhitelist.com/en/rationale.html>)
- From a sender's point of view, you will obviously want to get your IPv6 MTAs whitelisted at Spamhaus (and anywhere else doing whitelisting) as soon as possible.

Other Key Mail Technologies Are Already IPv6 Ready

- SPF supports IPv6 (see the “ip6 mechanism” at http://www.openspf.org/SPF_Record_Syntax)
- DKIM (and ADSP, for that matter) are IPv6 “agnostic” -- what they do doesn't depend on IPv4 or IPv6 addresses

So Now You're At Least "Dangerous..."

- You know what you need to do to get IPv6 connectivity, how to enable IPv6 on your servers, and how to enable IPv6 in your MTAs.
- If you complete those tasks for your servers, your MTAs will begin to send mail via IPv6 when the opportunity to do so presents itself.
- **Bail Out Opportunity #1:** If that's all you're looking for, you can now leave and find a suitable beverage -- you now know enough to at least be "dangerous." If, however, you're a glutton for punishment, you can stick around.
- Next we'll explain why others are finding IPv6 a matter of some urgency, and not an optional experiment (as it is for senders), and we'll also talk about some IPv6 "myths."

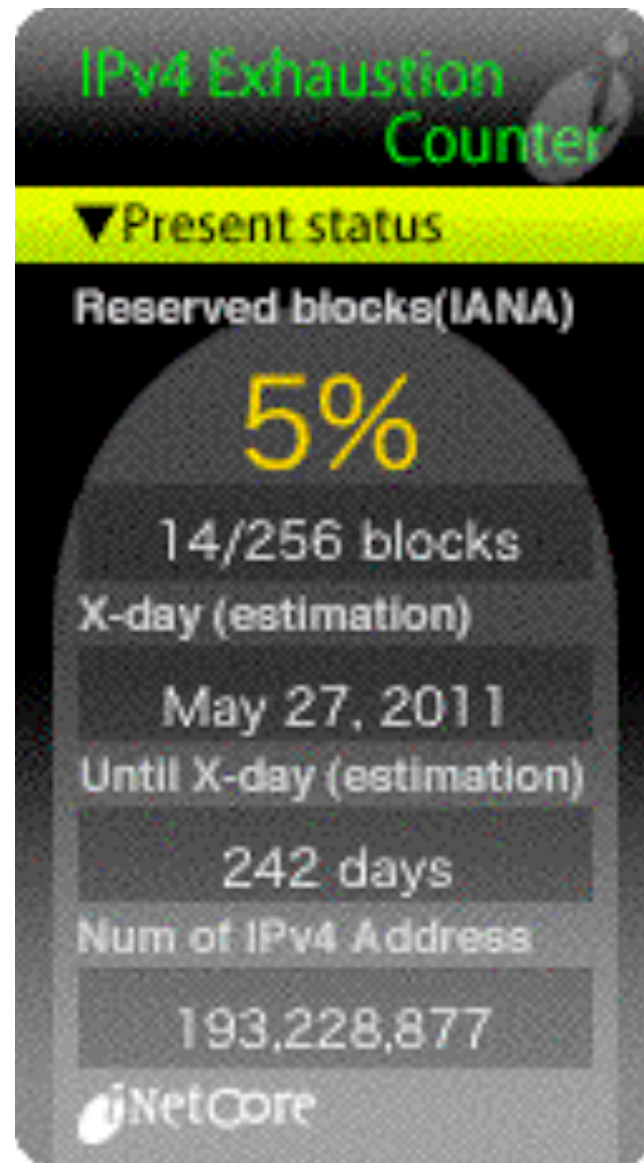
PART II.
**Understanding the Drivers and
Myths Behind IPv6**

6 (a). Understanding the Drivers for IPv6: IPv4 Address Exhaustion

IPv4 Addresses: A Scarce Resource

- There is a finite pool of available IPv4 addresses, and we're getting really, really close to running out.
- Based on the best available forecasts, see <http://www.potaroo.net/tools/ipv4/index.html> , the last IPv4 blocks will be allocated by IANA on 27-May-2011
- The regional internet registries (RIRs), such as ARIN, RIPE, APNIC, LACNIC and AFRINIC will exhaust the address space they've received from IANA less than a year later, around 21-Jan-2012
- These best estimates are based on current trends, and actual exhaustion might accelerate (or might slow) depending on what the community does (but probably not by much). From now till 21-Jan-2012 is roughly 1 year 3 and a half months away. That's not much time.

inetcore.com/project/ipv4ec/index_en.html



The Internet, Post-IPv4 Run Out

- Running out of IPv4 addresses isn't like running out of water in the middle of the desert, or air while SCUBA diving -- if you already have IPv4 address space, the space you already have will continue to work just fine.
- People who WILL run into problems, however, include:
 - growing ISPs who need more IPv4 IP addresses
 - new ISPs who need IPv4 addrs just to get started
 - customers of existing IPv4-based ISPs who may need to access network resources available ONLY via IPv6
 - customers behind weird/broken stopgap kludges
- Eventually, we risk the bifurcation of the Internet: part of the Internet may cling to IPv4 addressing, while the rest may end up having no choice but to use IPv6 addressing. Eventually, this will be a serious issue.

6 (b). Understanding the Drivers for IPv6: Regaining Internet Transparency

“But What About NAT?”

- While some sites (including uoregon.edu) assign each system on campus a globally routable IP address. Other sites (including many home users and many corporate sites) routinely employ network address translation (or “NAT”). NAT (actually PAT) makes it possible for multiple workstations to all use a single shared globally routable IPv4 address. If all you do is browse the web or use a web email service such as Hotmail, or Yahoo! Mail, or Gmail, NAT may superficially work fine for your needs.
- On the other hand, if you want to do Internet video conferencing, or use peer-to-peer applications, or you’re trying to track down and fix malware-infested hosts connecting from behind a NAT, you may find that NAT will make your life significantly more difficult.

NAT: A (Semi) Protocol-Aware Protocol

- Some network protocols (such as H.323) embed IP addresses in the traffic generated by those protocols.
- Because NAT rewrites network addresses, it needs to know HOW each protocol embeds IP addresses in network traffic streams. That is, NAT boxes need to be “protocol aware,” and thus networks using NAT are NOT “end-to-end transparent.” (Packets get rewritten during transport while passing through a NAT)
- If a NAT box faces traffic of a type that it doesn’t know how to handle, such as some new protocol, it can’t rewrite that traffic, and as a result that application will fail when run behind a NAT. This is very commonly the case for H.323 video conferencing, for example.
- Because of this, NAT’d networks can stifle application layer network innovation, or at least make it far harder!⁴⁰

The “Two Port Internet”

- Because of the problems that application developers face getting past NAT boxes (and restrictive firewalls!) it is common for developers to implement new protocols over http instead of developing new native protocols. Some of my colleagues refer to this as the “two port Internet” -- in this model, virtually all user traffic is either http (port 80) or https (port 443).
- Obviously this is something of an exaggeration (they forgot about DNS for example :-)), but it isn't entirely an argument w/o merit. All you need to do is look at network traffic and try to identify what applications make up most of the traffic to see the problem -- you can't do it any more just based on ports.
- C.F.: “A Look at the Unidentified Half of Netflow,” pages.uoregon.edu/joe/missing-half/missing-half.pdf⁴¹

End-To-End Transparency

- If you'd like to read about the importance of end-to-end transparency, some excellent starting points are:
 - RFC2775, "Internet Transparency," B. Carpenter, February 2000, <http://tools.ietf.org/rfc/rfc2775.txt>
 - RFC4924, "Reflections on Internet Transparency," B. Aboba and E. Davies, July 2007, <http://tools.ietf.org/rfc/rfc4924.txt>
- While Internet transparency is less often mentioned than imminent IPv4 address exhaustion as a reason why we need to deploy IPv6, transparency is nonetheless a very important underlying motivation for IPv6.

6 (c). Understanding the Drivers for IPv6: Controlling Route-Table Bloat


Controlling Route Table Bloat

- Another important (if little recognized) reason for promoting use of IPv6 has been the need to control the growth in the size of the global routing table. In fact, RFC4984 (<http://www.ietf.org/rfc/rfc4984.txt>) states,

"[...] routing scalability is the most important problem facing the Internet today and must be solved [...]"

The IPv4 Route Table Continues to Grow...

350,000



Source: <http://bgp.potaroo.net/as6447/>

IPv6 Was Supposed to Help Fix That

- When IPv6 was designed, address assignment was supposed to be hierarchical. That is, ISPs would be given large blocks of IPv6 address space, and they'd then use chunks of that space for each downstream customer, and only a single entry in the IPv6 routing table would be needed to cover ALL the space used by any given ISP and ALL their downstream customers (see RFC1887, "An Architecture for IPv6 Unicast Address Allocation")
- But now, let's pretend that my Internet connectivity is important to me, so I don't want to rely on just a single ISP -- I want to connect via multiple ISPs so that if one provider has problems, the other ISPs can still carry traffic for my site. This connection to multiple sites is known as "multihoming."

If I'm Multihomed, Whose Address Space Do I Use?

- When I get connectivity from sites A, B and C, whose address space would I announce? Address space from A? Address space from B? Address space from C? No...
 - A doesn't want me to announce part of its address space via B and C
 - B doesn't want me to announce part of its address space via A and C
 - C doesn't want me to announce part of its address space via A and B.
- I need to either assign each host multiple addresses (e.g., one address from A, one from B, and one from C), or I need to get my own independent address space which I can use for all three ISPs, but which will then take up a slot in the global routing table.

The Original Multiple IP Approach in IPv6


- The multiple IP approach was the original “answer” to this question in the IPv6 world.
- But if I assign multiple IPs to each host, one for each upstream ISP I connect to, how do I know which of those IP addresses I should use for outbound traffic generated by each host? Do I arbitrarily assign the address from A to some traffic? The address from B to other traffic? What about the address from C?
- Which of those addresses do I map to my web site or other servers via DNS? Do I use just A’s address? Just B’s? Just C’s? All three of those addresses? What if one of my providers goes down? Will traffic failover to just the other two providers quickly enough?

The Multihoming Reality Today

- IPv6 multihoming without use of provider independent address space is one of the unsolved/open issues in the IPv6 world today. Operationally, in the real world, ISP customers who need to multihome request their own provider independent IPv6 address space (cue Sonny and Cher: “The beat goes on, and the beat goes on...”)
- Route table growth may be a critical issue facing the Internet in the long term, but for now, the community has “dropped back into punt formation,” and we’re doing what needs to be done (at least for now) to get IPv6 deployed in a robust way (e.g., with multihoming). The good news is that the IPv6 table is still small (so we still have time to solve the IPv6 routing table growth issue); the bad news is that the IPv6 table is still small (which means many people still haven’t deployed IPv6!)

IPv6 Route Table Growth

3500



Source: <http://bgp.potaroo.net/v6/as6447/>

6 (d). Understanding the Drivers for IPv6: Regulatory Compliance

Federal Networks, For Example, Are Supposed to Be IPv6 Ready

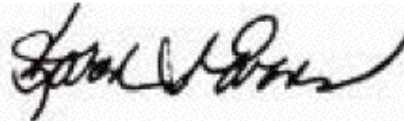
M-05-22

August 2, 2005

MEMORANDUM FOR THE CHIEF INFORMATION OFFICERS

FROM:

Karen S. Evans
Administrator



Office of E-Government and Information Technology

SUBJECT:

Transition Planning for Internet Protocol Version 6 (IPv6)

As I stated in my testimony of June 29, 2005, before the House Committee on Government Reform, we have set June 2008 as the date by which all agencies' infrastructure (network backbones) must be using IPv6 and agency networks must interface with this infrastructure. This memorandum and its attachments provide guidance to the agencies to ensure an orderly and secure transition from Internet Protocol Version 4 (IPv4) to Version 6 (IPv6). Since the Internet Protocol is core to an agency's IT infrastructure, beginning in February 2006 OMB will use the Enterprise Architecture

Source: www.whitehouse.gov/omb/memoranda/fy2005/m05-22.pdf

Theoretically All Federal Networks (At Least Temporarily) Met That Mandate, But...

- Reportedly many federal networks, having passed one IPv6 packet (and thus, however briefly, demonstrated that their backbones were IPv6 capable), promptly “re-disabled” IPv6 (ugh!)
- Check your favorite fed sites -- are they v6 accessible?
See: <http://www.mrp.net/cgi-bin/ipv6-status.cgi>

IPv6 Status Results

Domain Name	omb.gov	
HTTP	FAIL	No service could be found via IPv6
SMTP	FAIL	No service could be found via IPv6
DNS	0/0/8	No DNS service could be found via IPv6
NTP		No service defined
XMPP		No service defined

Let me know if you want this domain added to the survey page :-)
I'd also consider acting as a secondary DNS if you want an IPv6 accessible one.

© 2010 [Mark Prior](#)

IPv6 and Fed Scorecard Network Homepages?

www.dhs.gov --> no

www.doc.gov --> no

www.dod.gov --> no

www.doe.gov --> no

www.dot.gov --> no

www.ed.gov --> no

www.epa.gov --> no

www.hhs.gov --> no

www.hud.gov --> no

www.doi.gov --> no

www.doj.gov --> no

www.dol.gov --> no

www.nasa.gov --> no

www.nsf.gov --> no

www.nrc.gov --> no

www.opm.gov --> no

www.sba.gov --> no

www.ssa.gov --> no

www.state.gov --> no

www.usaid.gov --> no

www.usda.gov --> no

www.ustreas.gov --> no

www.va.gov --> no

Or pick another federal agency of your choice: the pattern is pretty consistent I'm afraid...

“Planning Guide/Roadmap Toward IPv6 Adoption Within the US Government”

- This is a new document (ca. May 2009) from the Federal CIO Council Architecture and Infrastructure Committee, Technology Infrastructure Subcommittee, Federal IPv6 Working Group, see <http://tinyurl.com/fed-cios-ipv6>
- I quote: “The purpose of this document is to provide U.S. government agency leaders with practical and actionable guidance on how to successfully integrate Internet Protocol version 6 (IPv6) throughout their enterprise. [...] without a concentrated effort by Federal agencies to effectively and efficiently deploy secure IPv6 network services, the Government’s technical advancement and ability to meet its mission needs will be critically impacted during the next 2 to 3 years.”

And Just This Last Week...

- On Sept. 28th, 2010, the NTIA held a workshop at which Federal CIO Vivek Kundra announced a directive “requiring all U.S. government agencies to upgrade their public-facing Web sites and services by Sept. 30, 2012, to support IPv6...” and that access must be via native IPv6 rather than an IPv6 transition mechanism.
- A second deadline, Sept. 30th, 2014, applies for federal agencies to upgrade internal client applications that communicate with public servers to use IPv6.
- For more, see
“White House Issues IPv6 Directive,”
<http://www.networkworld.com/news/2010/092810-white-house-ipv6-directive.html?page=1>

A Major Potential Stumbling Block: Non-IPv6 Content Delivery Networks (CDNs)

- Many federal web sites (and key commercial web sites) use Akamai (or another CDN) in order to handle huge online audiences, deliver good performance worldwide, and to resist DDoS attacks.
- For example, `www.irs.gov` is actually just a cname for `www.edgeredirector.irs.akadns.net`; whois confirms that `akadns.net` actually belongs to Akamai.
- If Akamai doesn't do IPv6, will current major Akamai customers (such as Apple, Cisco, Microsoft, RedHat, the Whitehouse, etc.) be willing to deploy IPv6 for critical sites without them?
- BTW, at least one vendor, Limelight, DOES offer an IPv4 and IPv6 CDN service...

But Speaking of Akamai, Akamai Is Reportedly Working On IPv6...

- I'm happy to report that Akamai is now reportedly working on IPv6-ifying its CDN infrastructure. See, for example, the coverage in:

"Akamai: Why Our IPv6 Upgrade Is Harder Than Google's,"

<http://www.networkworld.com/news/2010/091610-akamai-ipv6.html>

September 16th, 2010

The Issue Isn't Just Web CDNs...

- A growing number of sites outsource their email operations. Unfortunately some email-as-a-service (and some cloud-based spam filtering services) don't support IPv6, thereby limiting the ability of their customers to integrate IPv6 into their existing IPv4-based services.
- CDNs and outsourced email and spam filtering services aren't the only reason why IPv6 adoption has been slow at some major Internet sites, but they can certainly be stumbling blocks that will need to get resolved.
- Some of these outsourced email/spam filtering services and CDNs are being actively monitored on MRP's survey page, see www.mrp.net/IPv6_Survey.html , including Akamai, Google, McAfee, MessageLabs, MacQuarie, Microsoft, Postini, ProofPoint, RedCondor and WebSense.

7 (a). Understanding IPv6 Myths:
'IPv6 Will Improve "Network Security"
Due to IPv6 Having "Mandatory" IPSec'

(Sorry, No)

IPv6 and IPsec

- IPsec is not new with IPv6; in fact, IPsec dates to the early 1990's. What's different when it comes to IPv6 is that support for IPsec was made "mandatory" for IPv6 (see for example "Security Architecture for IP," RFC4301, December 2005 at section 10, and "IPv6 Node Requirements," RFC4294, April 2006 at section 8.)
- If actually used, IPsec has the potential to provide:
 - authentication
 - confidentiality
 - integrity, and
 - replay protection
- All great and wonderful security objectives -- IF IPsec gets used. Unfortunately, as we'll show you, what was supposed to be the cornerstone of the Internet's security architecture has proven in fact to be widely non-used.⁶¹

IPv6 Does NOT Require IPSec to Run...

- IPv6 can be brought up without IPSec getting enabled, and in fact this is routinely the case -- see an example on the next slide.
- More broadly, if people are doing cryptographically secured protocols of *any* sort, they inevitably run into problems -- crypto stuff just tends to be inherently complex and hard to learn to use. For example, how many of you routinely use PGP or GPG to cryptographically sign or encrypt your email, eh? How many of you are doing DNSSEC to cryptographically protect the integrity of your DNS traffic?
- Now think about how often you see people moaning about problems they're having getting IPSec to work with IPv6 -- do you EVER see that sort of thing on the mailing lists or discussion groups you're on? No, right? That's because hardly anyone is doing IPSec with IPv6.

Some IPv6 Traffic Statistics From A Mac OS X Host (Note That It's NOT Doing IPsec)

```
# netstat -s -finet6    <-- let's look just at IPv6 traffic  
[snip]
```

ip6:

```
124188 total packets received
```

```
[snip]
```

```
84577 packets sent from this host
```

```
[snip]
```

ipsec6: <-- how much of it was Ipsec protected?

```
0 inbound packets processed successfully
```

```
0 inbound packets violated process security policy
```

```
[snip]
```

```
0 outbound packets processed successfully
```

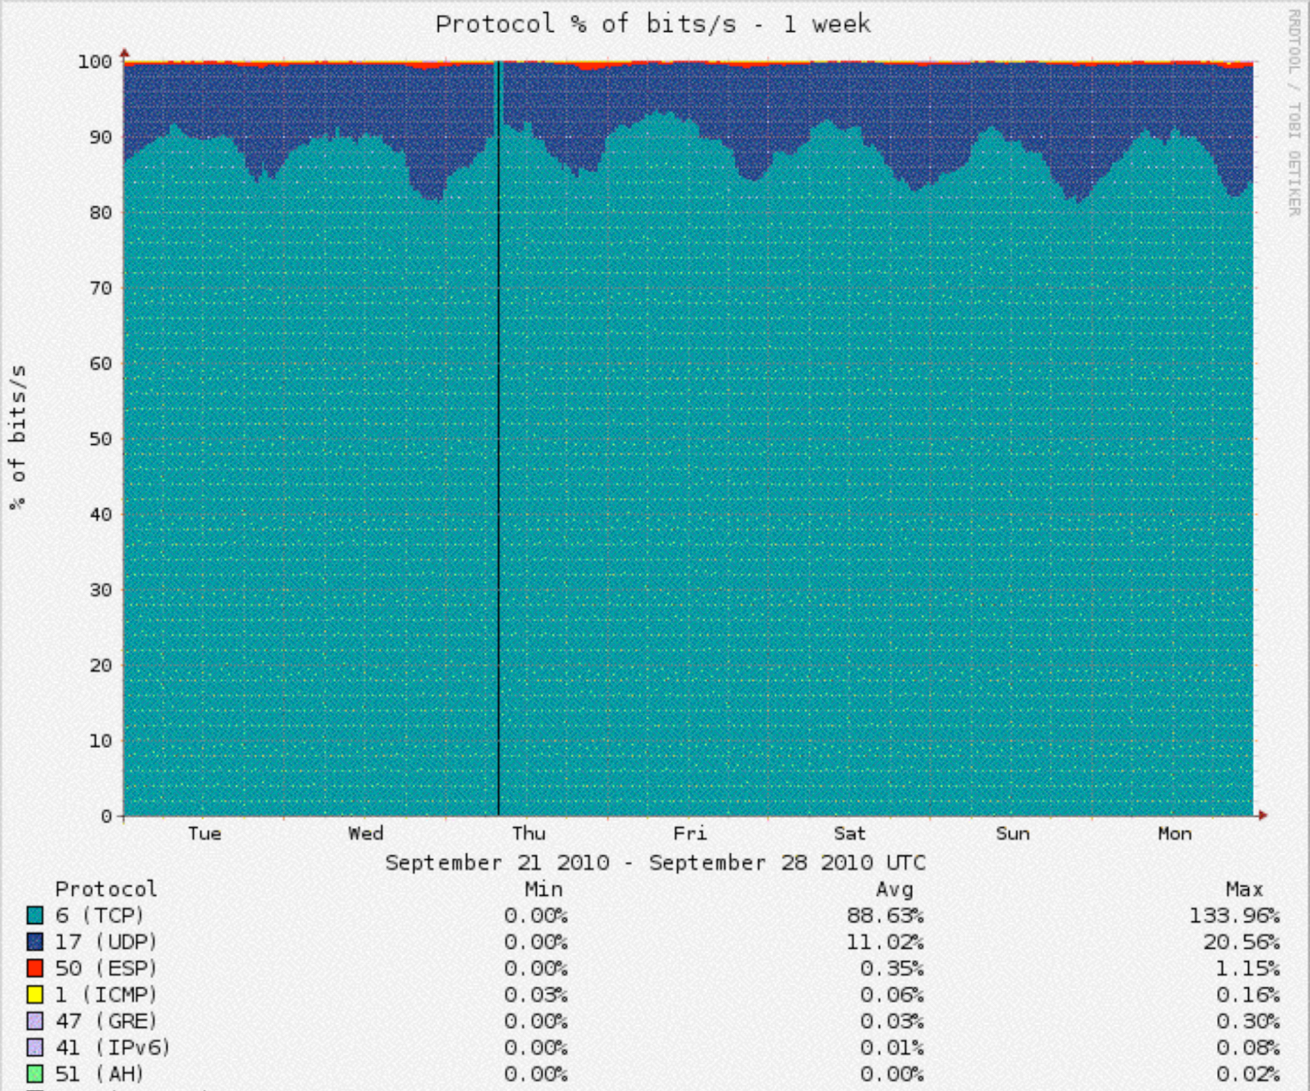
```
0 outbound packets violated process security policy
```

```
[snip]
```

IPsec (Even on IPv4!) Isn't Getting Much Use

- Raw IPsec traffic (AH+ESP, protocols 50 & 51) isn't seen much on the commercial IPv4 Internet.
- For example, about a year or so ago, Jose Nazario of Arbor Networks estimated IPsec traffic at 0.9% of octets (statistic courtesy the ATLAS project).
- CAIDA (thanks kc!) also has passive network monitoring data available; see <http://www.caida.org/data/passive/monitors/equinix-chicago.xml>
You can see the protocol distribution from a couple of CAIDA's monitors for one recent day on the next couple of slides. IPsec traffic is basically too small to even be seen for the most part.

Protocol Distribution From One of CAIDA's Passive Monitors



Bad News/Good News

- It would be foolhardy to expect IPsec to provide any material improvement to your site's security since the vast majority of your aggregate traffic (including virtually all your IPv4 traffic) will NOT be IPsec secured.
- On the other hand, the "good news" is that a lack of IPsec usage in the IPv6 world is substantively no worse than a lack of IPsec usage in the IPv4 world.

7. (b) Understanding IPv6 Myths:
“IPv6 Will Simplify Renumbering,
Improve Routing Performance By
Simplifying Packet Formats,
Improve Support For QoS,
Facilitate Mobility, etc.”

(maybe, but frankly, no one cares)

I Really Don't Mean to Sound Harsh, But...

- IPv6 may very well bring many cool new features to networking, but quite frankly, all of these incidental new features really don't matter -- they're not "make or break" drivers when it comes to adoption of IPv6.
- I'd love to hear hard evidence to the contrary, but truly, I've seen no indication that any of these other factors carry much weight in helping to shape the IPv6 go/no-go decision.
- On the other hand, there are some genuine reasons why people adamantly are NOT interested in doing IPv6.

**7. (c) Understanding IPv6 Myths:
Everyone's Running Out of IPv4
Address Space and Will Need IPv6**

**(No, Actually Some Sites Still Appear
To Have Quite A Bit of IPv4 Space Left)**

If You Have Abundant IPv4 Address Space, It May Be Hard to Get Excited About IPv6

- The definition of “abundant” will vary from site to site, but, for example, many universities have legacy /16’s, and 2^{16} (or roughly 65,500 addresses) can seem like “a lot” of addresses (even though they can go awfully fast when you have a campus of 20,000–25,000 people, most of whom have multiple networked devices, plus lots of printers and other networked infrastructure).
- Others, however, have /8’s (2^{24} , or roughly 16,777,200 addresses per /8), and that may be large enough to eliminate all worries at those sites about address scarcity.
- If you haven’t recently looked at the list of who has /8’s, you can check <http://www.iana.org/assignments/ipv4-address-space/ipv4-address-space>

Considering Just Large Federal Netblocks...

- The US government (typically as the US Department of Defense) controls a relatively large fraction of the entire Internet IPv4 space, including a dozen slash 8's: 6/8, 11/8, 21/8, 22/8, 26/8, 28/8, 29/8, 30/8, 33/8, 55/8, 214/8, and 215/8 (to say nothing of additional /8 blocks controlled by defense industrial contractors with close/extensive military contracts, plus miscellaneous smaller netblocks scattered hither and yon). Given that level of public IPv4 address space availability, one can understand that some federal agencies have not felt a particularly pressing need to move to IPv6 real ricky-ticky soon now.
- An aside about those large federal netblocks: even though addresses from some of those block may not show up in public routing tables, they are still being used, you're just not seeing them in some public routing tables.

Another Manifestation of the “Oh, No Worries, I’ve Got Plenty of Address Space” Phenomena

- While there are obviously only a comparative handful of sites which have their own public /8, many sites use RFC1918 “private” non-publicly routable address space, such as addresses from 10/8.
- If NAT currently meets your needs, as long as you have 10.0.0.0–10.255.255.255 (and 172.16.0.0–172.31.255.255, and 192.168.0.0–192.16.255.255) available, again, you may feel like your private IPv4 addressing needs are generally being well met (provided, of course, that you can also get the comparative handful of public IPv4 addresses you may also need).
- But, as we’ll see later, use of NAT forecloses some of the easiest IPv6 transition mechanisms, such as use of 6to4.

7 (d). Understanding IPv6 Myths:
'If You Enable IPv6, You'll Suddenly Have
Access to Lots of "Cool New Stuff" That's
Not Available via IPv4'

(Nope. Anything That's Available
Via IPv6, Will Also Be Equally
Available via IPv4)

An Important/Subtle Point to Understand

- Sometimes folks ask, “So if I begin to do IPv6, *what new stuff can I get at that I can’t get at already?*”
- Can you imagine Google, or Amazon, or Microsoft, or CNN, or <fill in the name of an important Internet site here> making their web site or other online resources ONLY available via IPv6? No, probably not.
- Any/all important Internet resources will ALWAYS be available via IPv4, even if those resources are ALSO available via IPv6.
- Thus moving to IPv6 does NOT magically give you access to new stuff that you couldn’t get to via IPv4 (well, technically there are a few IPv6-only things mentioned at <http://www.sixxs.net/misc/coolstuff/> , but nothing in and of itself that’s enough to justify deploying IPv6) 74

At Some Point In The Future, Though, The Default *Will* Change

- At some point in the future, people will eventually ask, “So if I still bother doing IPv4, *what sort of ‘old stuff’ can I get at that I can’t get at already via IPv6?*”
- We’re still a LONG way off from that point, but it WILL eventually happen.
- Remember when people used to carry AppleTalk or IPX or DECNet on their local area networks? They sure don’t anymore (or at least no one I know still does!)

7 (e). Understanding IPv6 Myths:
“All The Techies Will Figure Out Some New
Way To Stretch Out What IPv4 Space
We’ve Got Left -- I’m Not Going to Worry”

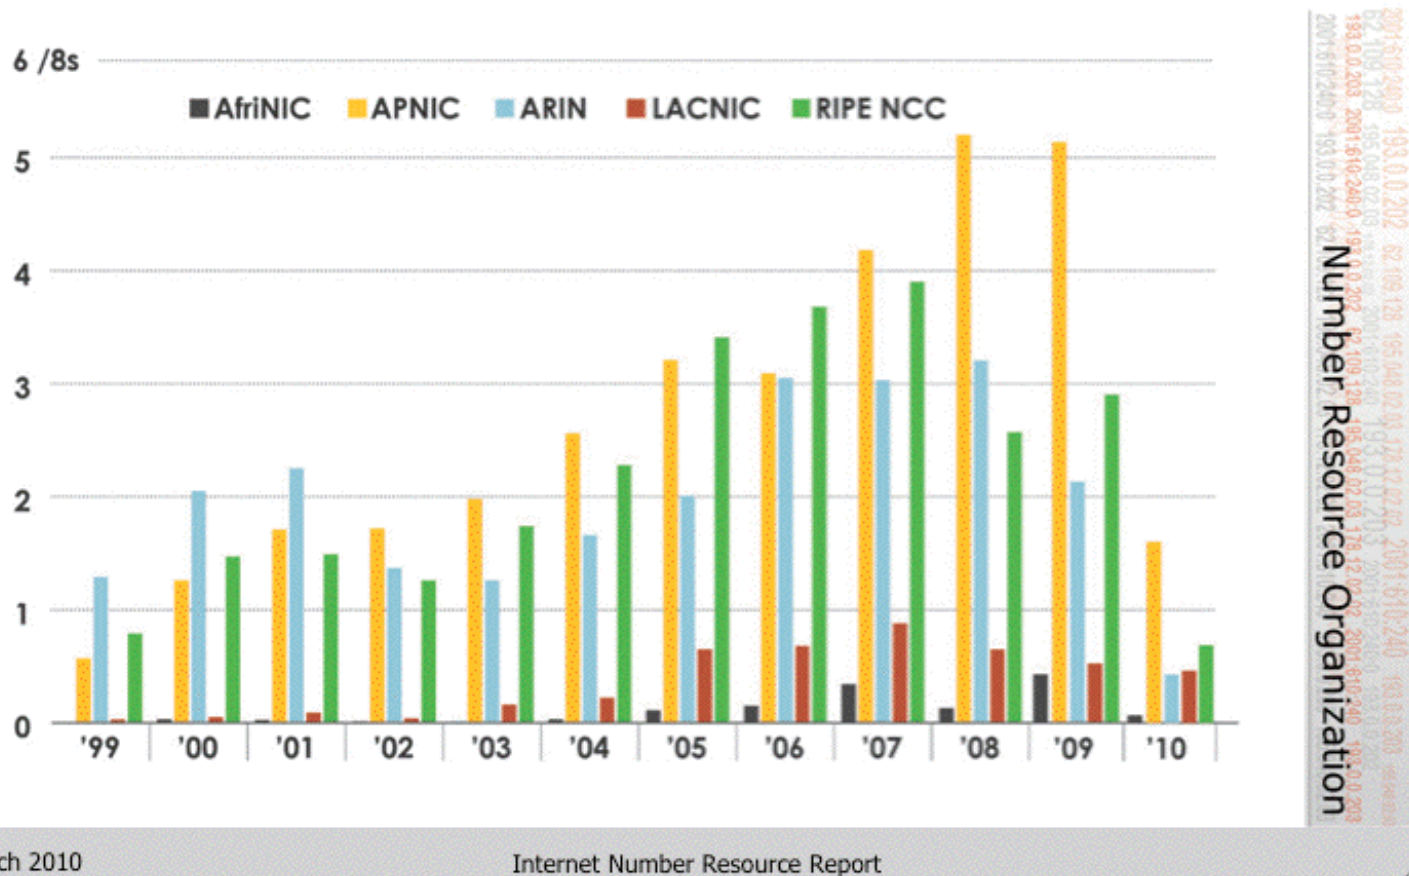
(I Admire Your Confidence,
But Personally, I Don’t Think
I’d Be So Sanguine)

We're Not The (Only) Ones Driving The Address Consumption Bus!



IPv4 ADDRESS SPACE ISSUED (RIRs TO CUSTOMERS)

In terms of /8s, how much space did each RIR issue by year?



March 2010

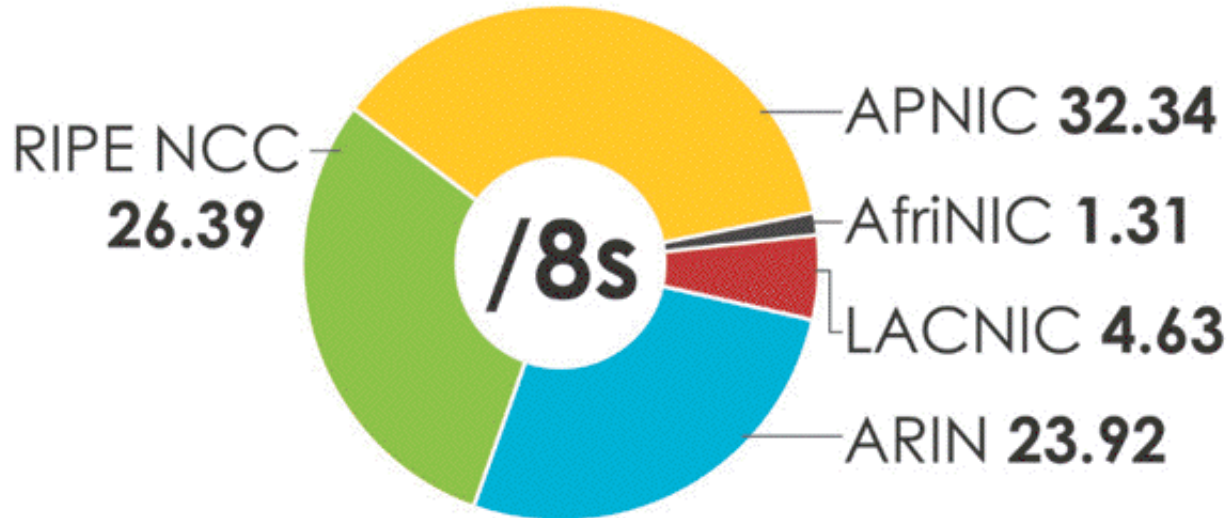
Internet Number Resource Report

A Cumulative View



IPv4 ADDRESS SPACE ISSUED (RIRs TO CUSTOMERS)

In terms of /8s, how much total space has each RIR issued?
(Jan 1999 – Mar 2010)



2001-6:10-2400 193.0.0.202 62 109 128 193.0.0.2 19 178 17 62 62 2001-6:10-240 193.0.0.203
193.0.0.203 2001-6:10-240-0 193.0.0.202 62 109 128 193.0.0.2 19 178 17 62 62 2001-6:10-240 193.0.0.203
2001-6:10-240-0 193.0.0.202 62 109 128 193.0.0.2 19 178 17 62 62 2001-6:10-240 193.0.0.203
Number Resource Organization

What If IPv4 Address Usage Was Proportionate to Regional Population?

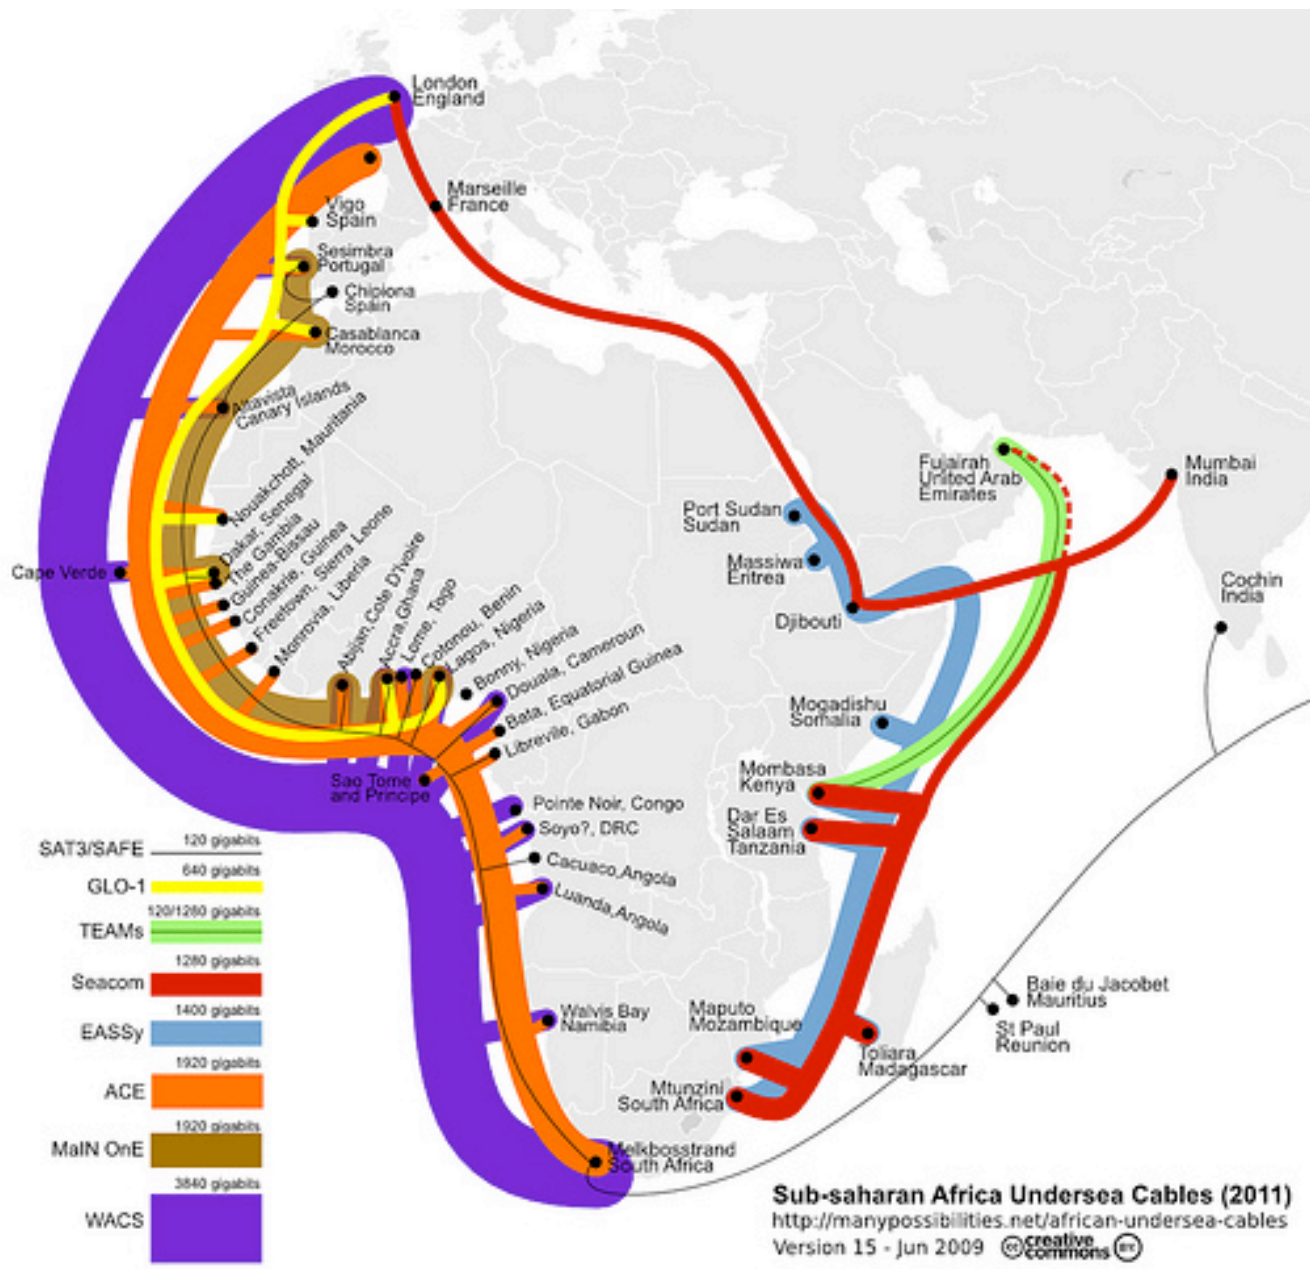
	Population	%	/8's	%	Ratio
● Asia:	4,121,097	60.3%	32.34	36.5%	0.605
● Africa:	1,009,893	14.7%	1.31	1.4%	0.095
● Europe:	732,206	10.7%	26.39	29.7%	2.775
● L. Amer.:	582,418	8.5%	4.63	5.2%	0.611
● N. Amer.:	348,360	5.1%	23.92	27%	5.29
● Oceania:	35,387	0.5%			
● Total:	6,829,360		88.56		

Population in thousands, mid year 2009 estimates

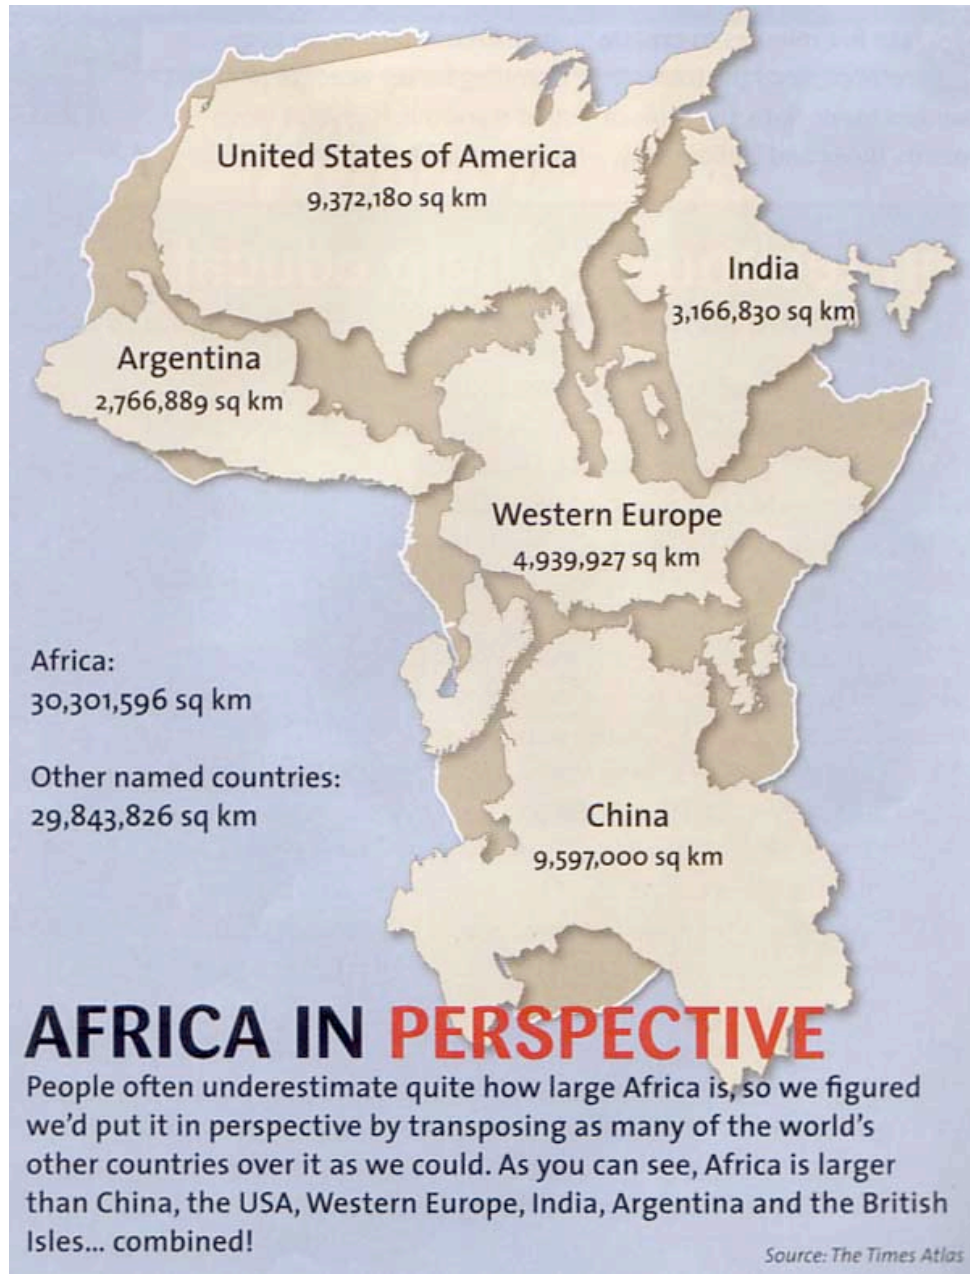
Note: Oceania's addresses are handled by APNIC (e.g., Asia)

The Future: Notice Relatively “Tiny” Africa

- Historically, Africa’s IP address usage to date has been minimal, less than one and a half /8s.
- This was likely due to a variety of factors, but at least one important factor was the high cost of connectivity (thousands of dollars per Mbps per month vs. just dollars per Mbps per month in the US (for large customers)).
- Another driver was widespread use of satellite Internet connectivity, with high latency, NAT’d connections and IPs provided by the satellite operator.
- Improved fiber connectivity is changing all that. Some of the world’s largest and most densely populated regions in both Africa and Asia are now coming online, and I believe the improved connectivity to those areas will result in a substantial demand for new IPv4 addresses.



http://blog.foreignpolicy.com/files/images/090618_africa_underseas_cables.jpg



http://strangemaps.files.wordpress.com/2006/11/africa_in_perspective_map.jpg

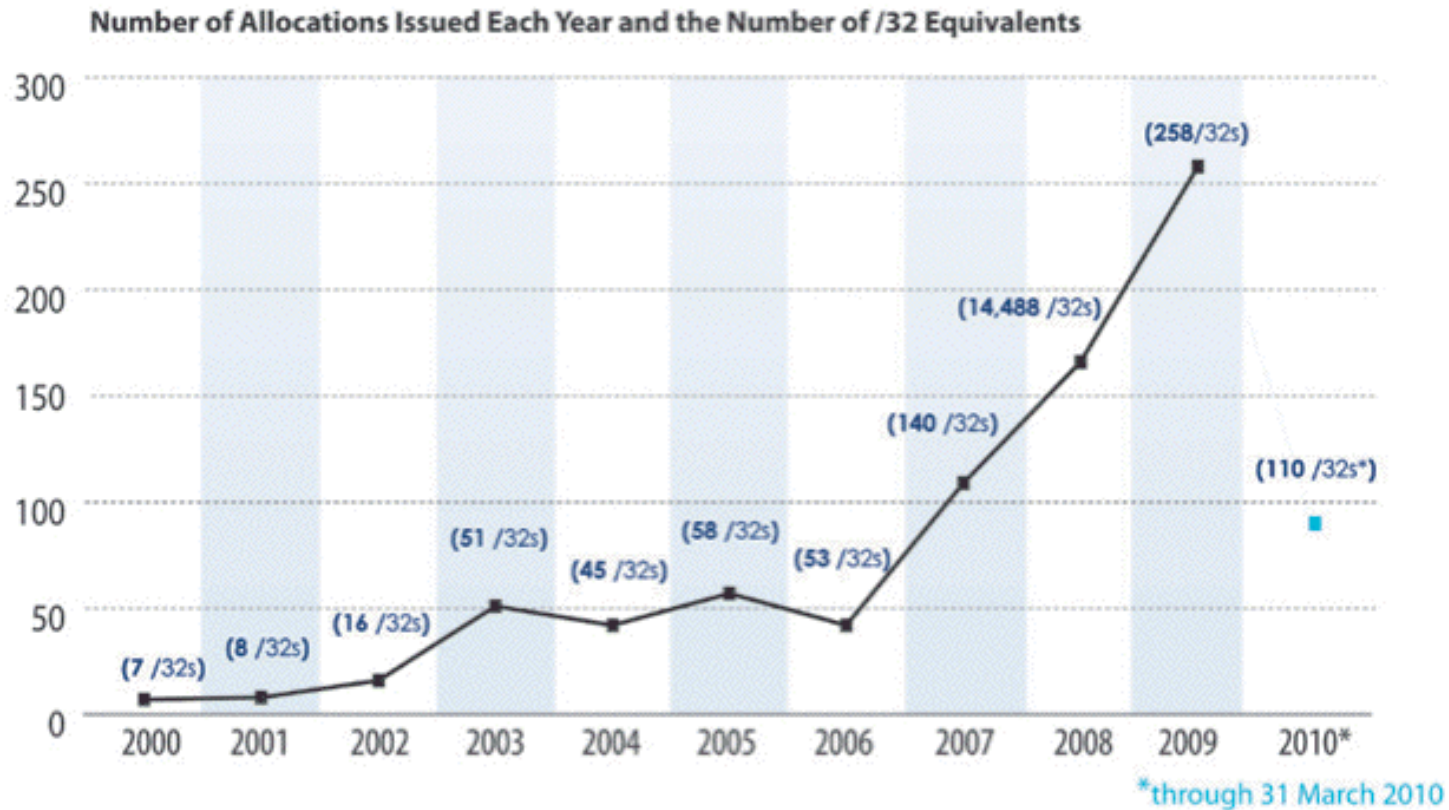
If You Believe We Have Enough IPv4 Addresses

- Given the preceding slides, you must also believe in miracles! :-)
- In my case, I'd rather be prepared to work with IPv4 AND IPv6 :-)

**7 (f). Understanding IPv6 Myths:
“Customers Just Aren’t Asking
For IPv6 (Except You!)”**

That's Not The Impression I'm Getting...

IPv6 Allocations Issued (/32s)



Source: https://www.arin.net/participate/meetings/reports/ARIN_XXV/PDF/Wednesday/Nobile_RSD.pdf

From the NANOG Mailing List, April 2010

- > What I heard at a recent (within the past six months)
 - > conference was that "there is no customer demand for v6" so it
 - > isn't on the immediate needs list. He said they had a lot of
 - > inquiries about v6, but to date not having native v6 wasn't a
 - > deal breaker with anyone

I watched a vendor at one conference tell 20 people in a row that each one of them was the only one asking for IPv6. I mentioned to him that he should have his short-term memory loss checked out by a physician. At first he was confused. When I pointed out what I had just seen him do, he went from confused to embarrassed and admitted that it was the party line from his marketing department and they knew IPv6 was important, but, didn't have a story to tell yet, so, they were trying to spin for delay.

<http://seclists.org/nanog/2010/Apr/71>


But The Perception Persists...

- Nobody actually cares*
 - Very little market demand
 - A “checkbox” feature
 - Few real endpoints

* Except Asia, US Government, Internet2

The Question That's Really Being Asked Is...

- "What's the business case?"
- Vendors who ask that question obviously fail to understand the technology adoption lifecycle:



<http://commons.wikimedia.org/wiki/File:Technology-Adoption-Lifecycle.png>

We're Crossing "The Chasm" Right Now

- The companies that **are** ready to support IPv6 will be well positioned to support emerging customer needs and to thrive in an IPv6 environment
- The companies that **are not** ready will find their customers doing what they need to do (nothing personal, it's just business, and since "no one but me" cares about IPv6, I'm sure you won't mind that I'm taking my business elsewhere, see ya...)

A 2nd Opportunity For You To Bail Out

- So that's the end of Part II of today's IPv6 discussion; the final chunk of our training today will begin to expose you to some technical material relating to IPv6.
- If you're not a "techie" and your head already hurts and you're not interested in diving in deeper, feel free to bail out now. You will not hurt my feelings if you do so. :-)

PART III.
IPv6 Technology Primer

8. (a): IPv6 Addresses

Starting With What We Know: IPv4

- IPv4 addresses are 32 bits long
- $2^{32}=4,294,967,296$
- Normally represented in “dotted decimal” format:
 - four 8 bit octets (0 to 255 decimal)
 - each octet is separated from the next by a dot
 - leading zeroes in each octet may be omitted
- Examples:
 - 127.0.0.1
 - 128.223.142.89
 - 64.170.98.32

Something A Little Different: IPv6

- IPv4 addresses are 128 bits / 32 hexadecimal digits long
- $2^{128}=340,282,366,920,938,463,463,374,607,431,768,211,456$
(e.g., 3.4×10^{38} addresses)
- Normally represented in "colon separated" format:
 - eight sets of four hex digits (0000 to FFFF hex)
 - chunks are separated with colons (:)
 - leading zeroes in each chunk may be omitted
 - for convenience, :: (two successive colons) may replace one or more all-zero chunks, but only once in any addr
- Examples:
 - 2001:48a8:6880:0095:0000:0000:0000:0021
 - ::1
 - 2001:468:d01:d6::80df:d617
 - fe80::203:93ff:febf:b638

Quick Quiz

- Structurally/superficially valid or invalid?
If invalid, why?
 - A) 2001:468:0d01:003c:0000:0000:80df:3c15
 - B) 2001:468:0d01:003c::80df:3c15
 - C) 2001:468:d01:3c::80df:3c15
 - D) 2001:760:2e01:1::dead:beef
 - E) 2001:480:10:1048:a00:20ff:fe9a:58c1:80
 - F) 2001:500::4:13::80
 - G) 2001:13G7:7002:4000::10
 - H) 2607:f278:4101:11:209:5bff:fe8f:6609
 - I) fe80::209:3dff:fe13:fcf7
 - J) ::

Quick Quiz Answers

- Structurally/superficially valid or invalid?
If invalid, why?
 - A) Valid
 - B) Valid
 - C) Valid
 - D) Valid
 - E) Invalid (nine chunks instead of eight)
 - F) Invalid (double colons appear more than once)
 - G) Invalid (G is not a valid hexadecimal digit)
 - H) Valid
 - I) Valid
 - J) Valid (albeit as the IPv6 “unspecified address”)

8. (b): IPv6 Prefixes

Starting With Something We Know: IPv4 Prefixes

- We originally had class A, class B and class C blocks:

Name	CIDR equiv	Number	Addresses Per Block
Class A	/8	128	16,777,216
Class B	/16	16,384	65,536
Class C	/24	2,097,152	256

- Like Goldilocks, some of those were too large, and some of those were too small.
- We needed something a little more flexible, and that's what we got from CIDR, or Classless Inter-Domain Routing.

Common IPv4 CIDR Prefix Lengths

- | | | | |
|---------|----------------------|---------|-----|
| /8 ==> | 16,777,216 addresses | /23 ==> | 512 |
| /9 ==> | 8,388,608 | /24 ==> | 256 |
| /10 ==> | 4,194,304 | /25 ==> | 128 |
| /11 ==> | 2,097,152 | /26 ==> | 64 |
| /12 ==> | 1,048,576 | /27 ==> | 32 |
| /13 ==> | 524,288 | /28 ==> | 16 |
| /14 ==> | 262,144 | /29 ==> | 8 |
| /15 ==> | 131,072 | /30 ==> | 4 |
| /16 ==> | 65,536 | /31 ==> | 2 |
| /17 ==> | 32,768 | /32 ==> | 1 |
| /18 ==> | 16,384 | | |
| /19 ==> | 8,192 | | |
| /20 ==> | 4,096 | | |
| /21 ==> | 2,048 | | |
| /22 ==> | 1,024 | | |

- IPv4 subnets are often /24's (or maybe /23's or /25's)

IPv6 Uses Prefixes, Too

- IPv6 prefixes are just a “little bit” bigger...

Prefix	IPv6 Addresses
/32	$2^{(128-32)} =$ 79,228,162,514,264,337,593,543,950,336 (that's like getting <u>4,294,967,296</u> /64 subnets)
/48	$2^{(128-48)} =$ 1,208,925,819,614,629,174,706,176 (that's like getting <u>65,536</u> /64 subnets)
/56	4,722,366,482,869,645,213,696 (that's like getting <u>256</u> /64 subnets)
/64	18,446,744,073,709,551,616 IPv6 addresses ¹⁰⁰ each

General Rules for IPv6 Allocations/Assignments/Subnetting

- Local Internet Registries (LIRs), e.g., ISPs, statewide networks, etc., will get one (or more) IPv6 /32 from ARIN, RIPE, APNIC, etc.
- Large sites will get an IPv6 /48 from their LIR's /32
- Small sites needing only a few subnets over 5 years will get an IPv6 /56 from their LIR's /32 (some LIRs will just always issue /48's)
- If one and only one subnet is needed, that entity gets an IPv6 /64 (some LIRs will just always issue /48's)
- Hosts get one or more IPv6 /128s

- ALL SUBNETS at most sites will normally be /64s (even if they have only one or a small handful of hosts).

- Do NOT try to get clever and do something exotic when it comes to subnetting in IPv6.

8. (c): Types of IPv6 Addresses

Types of IPv6 Addresses

- Global Unicast: 2000::/3
- Link Local Unicast: FE80::/1
- Loopback: ::1/128
- 6to4: 2002::/16
- Teredo: 2001:0000::/32
- Unique Local Unicast: FC00::/7
- Multicast: FF00::/8
- IPv4-Mapped: ::ffff:128.223.214.23
- Deprecated: Site Local addrs and IPv4-Compatible addrs.
- For more on IPv6 addresses, see RFC4291 (+ other RFCs).
- For the most part, we primarily care about global unicast, link local unicast, and loopback addresses

Address Type Discussion

- **Global Unicast** addresses are globally unique, “real” IP addresses. These are the way you’ll normally refer to most IPv6 hosts
- **Link Local Unicast** addresses are used for some purposes local to a particular link; outside the extent of that link, they don’t get used.
- **Loopback** addresses -- just like 127.0.0.1 in IPv4 space, the IPv6 loopback address (::1/128) is an internal virtual address that the server can use to refer to itself.
- **6to4 and Teredo** addresses are special IPv4-to-IPv6 transition mode technologies. As a sender, you won’t be using these.
- **Unique Local Unicast** addresses (RFC4193) are the IPv6 equivalent of RFC1918 IPv4 addresses. Don’t use them.
- **Multicast** addresses are just like multicast in IPv4, except that in IPv6 they’re used extensively on the LAN, but IPv6 multicast traffic rarely appear on the wide area Internet, unlike IPv4 where the exact opposite is largely true.
- **IPv4-Mapped** addresses allow hosts that only bind IPv6 sockets to also accept IPv4 addresses. (largely of interest to programmers)¹⁰⁴

Anycast

- IPv6 also uses "anycast" addresses. Anycast addresses look just like regular global unicast addresses, however the same prefix gets advertised from multiple locations.
- Through the magic of routing, users automatically connect to the closest anycast instance when accessing a UDP service (such as DNS) hosted on an anycast prefix.
- If one anycast instance goes down, routes adjust, and the other anycast instances seamlessly assume the additional load.

8. (d): Addresses and Systems

Expect Multiple Addresses

- When you look at the interfaces on an IPv6-enabled system, it will be routine for it to have multiple addresses/interface.
- Sometimes this will just be a single globally unique unicast address, plus a link local address, plus some other bits and pieces, other times you may see multiple globally unique unicast addresses.
- Do not let this shake you up.
- Speaking of multiple addresses, most hosts will have both IPv6 AND IPv4 addresses on some interfaces. This is known as being “dual stacked” and is perfectly normal and acceptable.

Looking at Addresses on Interfaces on Linux

```
% ifconfig -a
eth0      Link encap:Ethernet  HWaddr 00:09:3D:13:FC:F7
          inet addr:128.223.142.32  Bcast:128.223.143.255  Mask:255.255.254.0
          inet6 addr: 2001:468:d01:8e:209:3dff:fe13:fcf7/64 Scope:Global
          inet6 addr: fe80::209:3dff:fe13:fcf7/64 Scope:Link
          UP BROADCAST RUNNING MULTICAST  MTU:1500  Metric:1
          RX packets:1187468996 errors:0 dropped:1805 overruns:0 frame:0
          TX packets:1338373204 errors:0 dropped:0 overruns:0 carrier:0
          collisions:0 txqueuelen:1000
          RX bytes:232065679216 (216.1 GiB)  TX bytes:915094219311 (852.2 GiB)
          Interrupt:185

[snip]
lo        Link encap:Local Loopback
          inet addr:127.0.0.1  Mask:255.0.0.0
          inet6 addr: ::1/128 Scope:Host
          UP LOOPBACK RUNNING  MTU:16436  Metric:1
          RX packets:8143461 errors:0 dropped:0 overruns:0 frame:0
          TX packets:8143461 errors:0 dropped:0 overruns:0 carrier:0
          collisions:0 txqueuelen:0
          RX bytes:4295055907 (4.0 GiB)  TX bytes:4295055907 (4.0 GiB)
```

How Addresses Get On Interfaces: IPv4

- Starting with something you know, IPv4, most **end-user workstations** get their IP addresses automatically assigned via DHCP.
- Besides IPv4 addrs, RFC2132 describes additional useful bits that a host can get from a DHCP server, including:
 - subnet mask
 - router address(es)
 - domain name server(s)
 - the host name
 - the domain name
 - MTU information
 - broadcast address
 - plus an amazing amount of other bits and pieces (check it out if you don't believe me!)

How Addresses Get On Interfaces: IPv4 (2)

- Some sites may use host-specific DHCP entries to configure servers, too, but most IPv4 servers tend to be manually configured by a system administrator with a statically assigned IP addresses, host name and domain information, name server information, broadcast address, MTU information, etc.

How Addresses Get On Interfaces: IPv6

- In IPv6 world, most user workstations will get IPv6 addresses from *state-less address auto configuration*, or SLAAC. SLAAC runs on a router (not on a separate DHCP server), and as you might expect from the name, the IPv6 addresses that one gets via SLAAC are not maintained in a table anywhere (no “state” gets created when an IPv6 address is assigned via SLAAC).
- So how does the router know that it won't accidentally give you the same address as someone else (e.g., assign a duplicate address) if it doesn't keep track of who it has given an address to? Answer: it derives the address it gives you from something only you have, namely the MAC (hardware ethernet) address of your NIC

IPv6 Modified EUI-64 Format Identifier

- We need to take the 48 bit globally unique hardware MAC address from your network card, and convert that into a 64 bit dynamically assigned address. We do this by:
 - Taking the left most 24 bits of the MAC, and using that for the left most 24 bits of the EUI-64 ID
 - Taking the right most 24 bits of the MAC, and using that for the right most 24 bits of the EUI-64 ID
 - Cram the constant "FFFE" in the middle 16 bits
 - Tweak bit 7 (counting from the left) from zero to one (this is the "universal/local" bit)

The "front half" of the 128 bit address comes from the network (remember our rule that all subnets are /64s!)

“But Joe!” [Pretend You’re A Geek]

- “The whole world will know my unique and unvarying hardware MAC address! Evil marketers will track and correlate my every move wherever I may connect my IPv6 device, based on my invariant MAC address! This is worse than cookies!”
- True.
- This concern spawned another type of IPv6 address, so-called RFC3041 “Privacy Addresses.” These addresses effectively use a random address for the low order 64 bits of the IPv6 address, instead of a value derived from the host’s MAC address. Systems may also periodically change those addresses from time to time.

“BUT JOE!!!!”

- “That’s NOT what I wanted! If we give users random network addresses, how will we be able to track down abusers??? The same user may have one IPv6 address now, and another completely different IPv6 address later, and I don’t see how we’d keep track of who’s got what address when!!! What a pain!”
- One bit of potential happiness: NDPmon (think of this as “arpwatch for IPv6”). See ndpmon.sourceforge.net
- Another potential option: control network access just as you currently control network access for wireless networks (e.g., use 802.1x or a homegrown authentication solution)

Selecting/Deselecting Privacy Addresses

- **Windows:** privacy addresses are **enabled** by default when IPv6 is enabled on Windows XP. To disable them, see the next slide.
- **Macs:** privacy addresses are **disabled** by default. To enable them:
`# sysctl net.inet6.ip6.use_tempaddr=1`
- **Linux:** like Macs, privacy addresses are disabled by default. To enable them:
`# sysctl net.ipv6.conf.all.use_tempaddr=2`
`# sysctl net.ipv6.conf.default.use_tempaddr=2`
- Periodically recheck your assigned addresses if this is a big deal for you, and remember, this is NOT life-and-death privacy, it is just “something-to-make-life-hard(er)-for-intrusive-marketers”-grade privacy.

Disabling IPv6 Privacy Addresses



Disabling privacy addresses

- Windows XP

```
ipv6 -p gpu UseTemporaryAddresses no
```

- Windows 2003

```
netsh interface ipv6 set privacy state=disabled store=persistent
```

- Windows Vista

```
netsh interface ipv6 set privacy state=disabled store=persistent  
netsh interface ipv6 set global randomizeidentifiers=disabled  
netsh interface ipv6 set global randomizeidentifiers=disabled store=persistent
```

- Windows 2008

```
netsh interface ipv6 set global randomizeidentifiers=disabled  
netsh interface ipv6 set global randomizeidentifiers=disabled store=persistent
```

A Stateful Alternative to SLAAC: DHCPv6

- An alternative to SLAAC for workstations is DHCPv6, much like DHCP for IPv4.
- One critical difference: while DHCPv6 is well supported by IPv6ified versions of Microsoft Windows, at least some important vendors (read that as “Apple”) do not support DHCPv6 at this time (and may never do so).
- If you’re using a Mac OS X box and you want or need to do DHCPv6, you will need to run a third party DHCPv6 client (Dibbler is a commonly suggested option, see <http://klub.com.pl/dhcpv6/> , but note the comment there that “Due to work on my Ph.D, the Dibbler project is in the maintenance mode. Active development and non-critical bug fixing is on hold, until I finish my dissertation. Sorry.”)

8. (e): IPv6 DNS

If You Run Your Own Name Servers: Make Those Name Servers IPv6 Aware

- You'll need to think about making your name servers IPv6 aware. By this I mean your recursive resolvers and your authoritative name servers should both "know about" AAAA records.
- Your *access* to your name servers need not change -- you can continue to access them over IPv4 -- you just need to make sure that they understand AAAA records.
- Recent versions of BIND from ISC, the most common name server, are fully able to support IPv6
- If you have off site secondary name servers, make sure that those secondary name servers are also IPv6 aware
- A new option for folks is to outsource recursive DNS to Google -- 8.8.8.8 and 8.8.4.4 are already IPv6 aware¹¹⁹

IPv6 DNS Is Fairly Similar to IPv4 DNS

- In IPv4 world, servers and other hosts use “A” records to map fully qualified domain names to IP addresses:

```
% dig www.uoregon.edu a +short  
128.223.142.89
```

- In IPv6 world, we use “AAAA” (“quad A”) records instead of A records to map fully qualified domain names to IPv6 addresses:

```
% dig network-services.uoregon.edu aaaa +short  
2001:468:d01:3c::80df:3c15
```


Inverse Address Records Are Also Similar

- IPv4 world:

```
% dig -x 128.223.142.89 +short  
www.uoregon.edu.
```

- IPv6 world:

```
% dig -x 2001:468:d01:3c::80df:3c15 +short  
network-services.uoregon.edu.
```

- If you need a web-accessible IPv6 dig interface, try <http://www.digwebinterface.com/>

Complication: IPv6 Inverse Address Records

- There aren't a lot of good solutions for IPv6 inverse address record creation (other than manually creating them for servers, publicly visible network links, etc.).
- Dynamic DNS isn't such a hot idea, nor is building an inverse address record for millions of IPv6 addresses which may never get used!
- Ron Broersma from DREN described one potential approach in a talk he did at Joint Techs, see:
<http://www.internet2.edu/presentations/jt2010feb/20100202-broersma.pdf>
but it is unclear if the code he mentions there will be available for others to also use. (I hope it will be!)

Complications: IPv6 AND IPv4 Domain Names

- If a fully qualified domain name is bound to both IPv4 and IPv6 addresses, which one gets used? Which one should be “preferred?” The IPv6 one or the IPv4 one?
- This may be determined by the application (e.g., it may ask for both, and then use its own internal precedence information to determine which it will use), or by the DNS server (if it can give you an IPv6 address for a host, it will, and it will then stop; this is a problem if you advertise an IPv6 address for a host but then don't actually offer IPv6 connectivity for that domain name!)
- One sort of kludgy solution: put the IPv6 address in a subdomain. E.g., `www.example.com` (IPv4 only) might also have a related domain `www.ipv6.example.com` (IPv6 only)

What Address Do I Appear to Be Using?

- Sometimes you just want to know what address your system appears to be using.
- When that's the case, try:

<http://whatismyv6.com/>

- If you find yourself routinely curious about whether you're connecting to a site via IPv4 or IPv6, you can install the ShowIP Firefox addon available from

<https://addons.mozilla.org/en-US/firefox/addon/590>

It will then report the IP address that it is using for each web site you visit, green for IPv6 and red for IPv4

Protocol Precedence More Generally

- Assume you have IPv4 connectivity as well as IPv6.
- If a host is available via both IPv4 and IPv6, what's the "right" order (or "precedence") with which traffic to that dual stack host should be tried? Should traffic go over IPv4 whenever possible, or over IPv6 whenever possible? Obviously IPv4 is the old-and-proven familiar/safe/production-quality choice, but the object of this whole exercise is to eventually move to IPv6, right?
- On the other hand, some IPv6 connectivity may be less direct/more circuitous than production IPv4 connections, with greater latency and maybe even some packet loss. Do we really want to force as much traffic as possible over that sort of "disadvantaged" connection?
- Thoughts?

8. (f): Enabling IPv6 in Desktop Operating Systems

Most Desktop OS's Will Autoconfigure (Except XP)

To enable IPv6, begin by putting up a Windows command prompt terminal window

Start ==> Programs ==> Accessories ==> Command Prompt

C:\> netsh interface ipv6 install (if you ever want to remove, change install to uninstall)

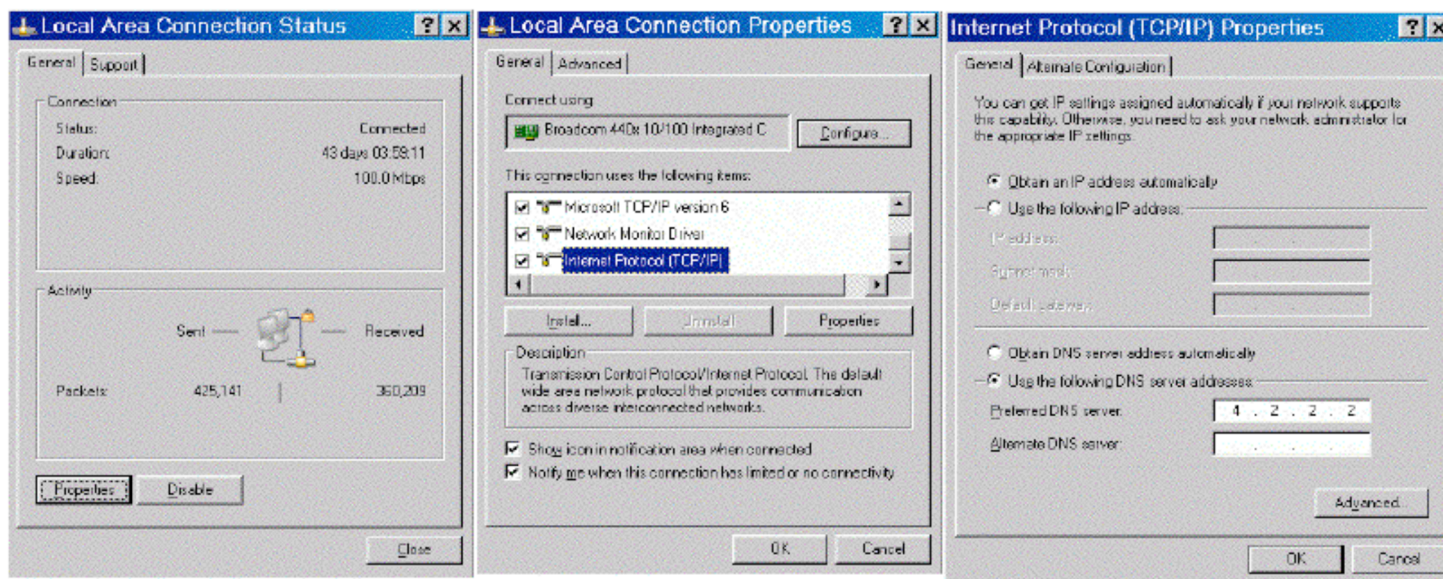
C:\> netsh interface ipv6 set teredo client (member of a domain? use enterpriseclient not client) (to disable, use disabled, not client)

C:\> ping ipv6.google.com (no reply? you probably need to configure an alternative IPv6 aware name server)

Configuring An IPv6 Aware Name Server, If Needed

Windows XP Service Pack 2 needs to one or more IPv6-aware name servers defined (an IPv6-aware name server is one which can return AAAA records for IPv6-connected fully qualified domain names). Note that Windows XP does **not** support name servers **accessible** only via IPv6. You may already have received a suitable name server as part of your existing IPv4 configuration, but if not, you will want to add an IPv6-aware name server. To manually configure an IPv6-aware name server, do:

Start ==> Settings ==> Control Panel ==> Network Connections ==> Local Area Connection



Click Properties.

Come down to Internet Protocol (TCP/IP) and select that item. Now click Properties.

Click Use the following DNS server address and enter the name of the IPv6 aware name server you want to use (such as 4.2.2.2).

Enabling the Windows Software Firewall for Teredo

By default, Teredo is not subject to the Windows software firewall that ships with Windows. To enable/confirm IPv6 firewall coverage:

Start ==> Settings ==> Control Panel ==> Windows Firewall ==> Advanced tab

Local Area Connection should be checked (and should stay checked)

Select Local Area Connection, then click on Settings

If Teredo is listed as an accessible service (e.g., it is checked), uncheck it

Click OK, OK

If you are using a 3rd party Windows software firewall instead of the integrated Microsoft Firewall, consult your vendor for advice with respect to using that firewall in a Teredo-enabled IPv6 environment.

Confirming You Have IPv6 Connectivity

To confirm that you have IPv6 connectivity, try tracing to an IPv6 only destination such as ipv6.google.com

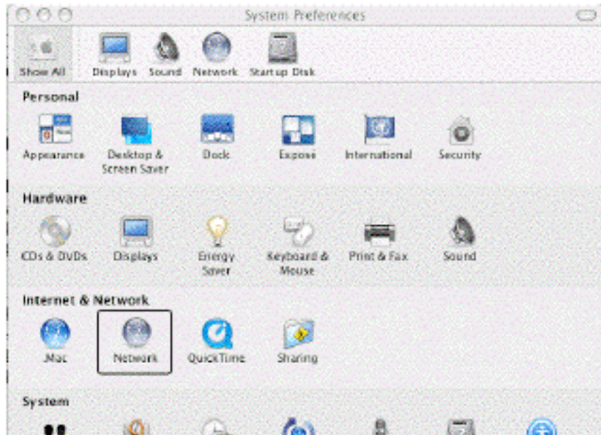
Start ==> Programs ==> Accessories ==> Command Prompt

```
C:\> tracert6 ipv6.google.com
```

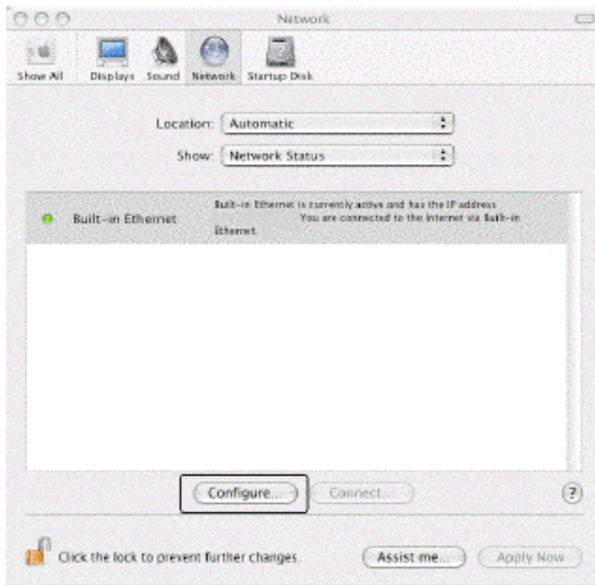
If you see a series of hops with IPv6 addresses, you have IPv6 connectivity. You may also find it interesting to see the path your traffic takes when using Teredo for IPv6 connectivity.

<http://www.uoregon.edu/~joe/ipv6/windows-xp-ipv6-with-teredo.pdf>

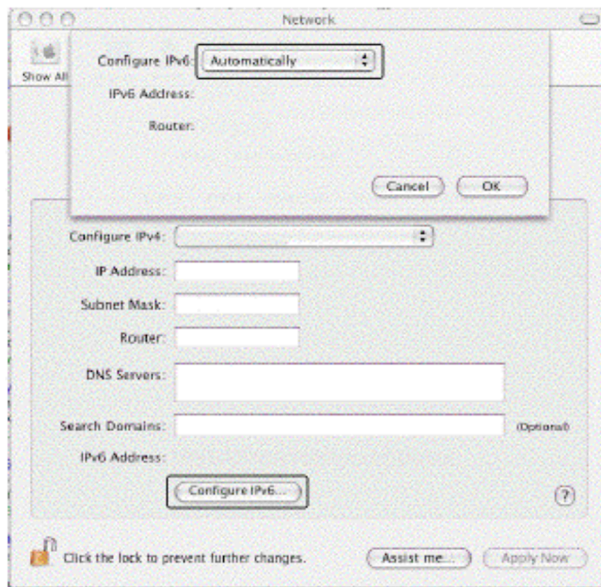
Mac OS X



Apple Menu ==> System Preferences ==> Network



On the network preferences screen, choose the drop down item: Network Status
Select the interface to be configured (e.g., wireless, built-in ethernet, etc.)
Click Configure



On the TCP/IP tab

Click on Configure IPv6 near the bottom of the pane

When the Configure IPv6 panel comes up, select Automatically

Click OK to close the Configure IPv6 panel

Click Apply Now

Close the Network Preferences window by clicking the red X button in the upper left corner

You should now have IPv6 enabled on your Mac

8. (g): IPv6 Web Browsers

Enabling IPv6 in Firefox for the Web

- If you don't already have Firefox, download it from:
<http://www.mozilla.com/firefox/>
- Firefox is IPv6 aware, but may need to be tweaked the first time you use it (it will remember this setting after that):
 - Launch Firefox
 - In the address bar, enter `about:config` and hit return
 - Filter on `ipv6`
 - Make sure `network.dns.disableIPv6` is set to `false`
 - If it isn't, double click on it to toggle it to `false`
- Now try going to <http://ipv6.google.com/>

A Somewhat Random Selection of IPv6-Accessible Web Sites

Google	http://ipv6.google.com/	2001:4860:b005::68
AFRINIC	http://www.afrinic.net/	2001:42d0::200:80:1
APNIC	http://www.apnic.net/	2001:dc0:2001:b:4608::115
Arctic Regional Supercomputer Center	http://www.arsc.edu/	2001:480:150:75::109
Argonne National Lab Public Software Mirror	http://mirror.anl.gov/	2620::dc0:1800:214:4fff:fe7d:1b9
ARIN	http://www.arin.net/	2001:500:4:13::80 and 2001:500:4:13::81
Berkeley	http://www.ipv6.berkeley.edu/	2607:f140:ffff:fff::80
DANTE	http://www.dante.net/	2001:798:2:284d::60
deepspace6.net (Linux IPv6 portal)	http://www.deepspace6.net/	2001:760:2e01:1::dead:beef
DREN	http://www.ipv6.dren.net/	2001:480:10:1050::5
ESNet	http://www.es.net/	2001:400:14:3::d
Fort Scott Community College	http://www.fortscott.edu/	2001:49d0:2c02::214:22ff:fe0f:38f4
Hexago	http://www.hexago.com/	2001:5c0:1000:10::2
HPCMO	http://www.hpcmo.hpc.mil/	2001:480:430:dddd:68b0:bff:fed5:f128
Hurricane Electric	http://ipv6.he.net/	2001:470:0:64::2
IANA	http://www.iana.org/	2620:0:2d0:1::193
ICANN	http://www.icann.org/	2620:0:2d0:1::103
IETF	http://www.ietf.org/	2001:1890:1112:1::20
Internet2	http://www.ipv6.internet2.edu/	2001:468:1420::151
Iowa State University	http://www.iastate.edu/	2610:130:101:100::7
IPv6.org	http://www.ipv6.org/	2001:6b0:1:ea:202:a5ff:fe0d:13a6
IPv6 Portal	http://www.ipv6tf.org/	2a01:48:1::2e0:81ff:fe05:4658
IPv6 Task Force	http://www.ipv6.eu/	2001:690:1fff:200::226
ISC	http://www.isc.org/	2001:4f8:0:2::d
KAME Project	http://www.kame.net/	2001:200::8002:203:47ff:fea5:3085
KANREN	http://www.kanren.net/	2001:49d0:3c00:1:209:6bff:fe7f:6c06
Kenya NIC	http://www.kenic.or.ke/	2001:43f8:10:50::2655
LACNIC	http://www.lacnic.net/	2001:13c7:7002:4000::10

Linkopings University	http://www.liu.se/	2001:6b0:17:f005::148
MAGPI	http://www.magpi.net/	2001:468:1802:101::805b:22c
Michigan Tech	http://www.ipv6.mtu.edu/	2001:48a8:0:2::1:22
NANOG	http://www.nanog.org/	2001:48a8:6880:95::21
Nederlandse IPv6 Task Force	http://www.ipv6-taskforce.nl/	2001:610:512::1000
NREN	http://www.nren.nasa.gov/	2001:4d0:8102:8:198:10:138:131
NTT	http://www.nttv6.net/	2001:fa8::80
Peter Bieringer	http://www.ipv6.bieringer.de/	2001:a60:9002:1::186:6
RIPE	http://www.ripe.net/	2001:610:240:11::c100:1319
Sauk Valley Community College	http://www.svcc.edu/	2001:470:c10b::214:5eff:fe28:7878
Space and Naval Warfare Systems Command	http://www.spawar.navy.mil/	2001:480:10:1048:a00:20ff:fe9a:58c1
SURFNet	http://www.ipv6.surfnet.nl/	2001:610:510::192:42:113:60
SWITCH	http://www.switch.ch/	2001:620:0:1b::b
TERENA	http://www.terena.org/	2001:610:148:dead::6
Tsinghua University	http://www.tsinghua.edu.cn/	2001:da8:200:200::4:100
UCLA	http://www.ucla.edu/	2607:f010:3fe:301:101d:9ff:fe32:a7db (etc)
University of Hawaii	http://www.ipv6.hawaii.edu/	2607:f278:4101:11:209:5bff:fe8f:6609
University of Oregon video archive	http://limestone.uoregon.edu/	2001:468:d01:103::80df:9d0c
Virginia Tech	http://www.ipv6.vt.edu/	2001:468:c80:210f::173:2bc6:4145
WPI	http://www.wpi.edu/	2001:468:616:824::31
3ROX	http://www.3rox.net/	2001:5e8:0:1000::fa
6TAP	http://www.6tap.net/	2001:400:14:3::9

Track the IPv6-ification of some of your favorite web sites at http://www.mrp.net/IPv6_Survey.html

Access most any IPv4-only web site via IPv6 using sixxs.org's IPv4-to-IPv6 gateway by just appending sixxs.org to the normal IPv4 only domain name. For example, to access www.cnn.com via IPv6, you'd go to <http://www.cnn.com.sixxs.org/>

What About Other GUI Web Browsers?

- Opera (<http://www.opera.com/>) is ready to use IPv6 as shipped with no tweaking required, likewise Internet Explorer 8.
- Safari is IPv6 enabled, but some versions of Safari may prefer IPv4 to IPv6 when both are available. To reverse that preference for dual stack sites:
 - 1) In a Mac terminal window (Applications --> Utilities --> Terminal) enable the Safari Debug menu by entering:
defaults write com.apple.Safari IncludeDebugMenu 1
You can then close the terminal window.
 - 2) Launch Safari, then go to Safari's Debug Menu (on the far right side of the main Safari menu), and come all the way down to Supported Protocols and uncheck http: (Simple Loader). You only need to do this once.

8. (h): IPv6 Email Clients

Modern Mail User Agents Support IPv6

- Apple Mail.App
- Opera Mail client
- Outlook 2007
(see <http://support.microsoft.com/kb/924469>)
- Thunderbird: (see the next slide)
- Other MUAs?

Enabling IPv6 in Thunderbird for Email

- Thunderbird (Mozilla's graphical email client, sibling application to the Firefox web browser) may ship with IPv6 DNS queries disabled by default.
- Enabling IPv6 DNS queries requires you to:
 - Start Thunderbird
 - Go to Thunderbird --> Preferences
 - Go to Advanced (the "sprocket" or "gear" at the top)
 - Click the "Config Editor" button near the bottom
 - In the about:config window's Filter: box, enter `ipv6`
 - You should see:
`network.dns.disableIPv6 [etc] false`
 - If, however, that setting is true, double click true to toggle it to false

Public IPv6-Accessible Web Email

- As another example of IPv6 email lagging relative to where it should be, consider web-based email, one of the most popular and widely used applications on the Internet today.
- To the best of my knowledge, there is currently only ONE (1) public web email provider which has even an experimental web email service accessible via IPv6, and that's <http://ipv6.rollernet.us/>
- Technically, if you're able to manipulate your local hosts file, you can also access Google's Gmail via IPv6, see <http://jeremy.visser.name/2008/11/25/how-to-access-gmail-and-google-reader-over-ipv6/> (URL wrapped due to length), but that's hardly the sort of thing that a typical user should be expected to be able to do.
- If you're aware of any additional IPv6-accessible public web email service providers, I'd love to know about them.

8. (i): ssh for IPv6

ssh and IPv6

- ssh supports IPv6. To get OpenSSH (5.5p1 was released 4/15/2010!) go to www.openssh.com/portable.html (be sure you're also running the latest OpenSSL -- 1.0.0 was released 3/29/2010 -- see www.openssl.org/ and don't forget to set a sane installation prefix!)
- To force ssh to use IPv6, use:
`% ssh -6 joe@example.com`
- To force ssh to use IPv4, use:
`% ssh -4 joe@example.com`

8. (j): Web Servers

What About Your Web Site?

- Is your web site IPv6 accessible? If not, why not? The most popular web servers, e.g., Apache and Microsoft IIS, both support IPv6 these days! So what's the "holdup," eh?
- One potential roadblock to explore is whether hardware load balancers are IPv4 only. Hardware load balancers commonly sit in front of multiple physical computers, making a pile of systems act as if they were one computer. Some popular hardware load balancers know how to do IPv6, but license IPv6 functionality separately from IPv4.
- Other times the issue may be the use of outsourced content delivery networks (CDNs) which may not be IPv6-enabled (previously discussed)
- Lastly, at least sometimes, no one may every have asked, "Hey, why isn't our web site IPv6 accessible?" Perhaps that's a question that YOU should ask once you're done with this training?

Don't Forget About IPv6 Adrs in Log Files

```
# cd /usr/local/apache2/logs
# cat access_log
[...]
2001:468:d01:d6::80df:d617 - - [23/Apr/2010:10:20:29 -0700]
"GET / HTTP/1.1" 200 54
[etc]
```


Does *your* log file analyzer product support IPv6 addresses? Some, like AWStats from <http://awstats.sourceforge.net/> require a separate plugin to enable some IPv6 functionality; other functionality, like mapping addresses to geographic locations, may simply not be available for IPv6

Conclusion

Learning More About IPv6: Books

- A nice online reference: 6Net IPv6 Cookbook, www.6net.org/publications/deliverables/D3.1.2v2.pdf
- Traditional printed books about IPv6 are also available. If you go to Amazon's book section and search for IPv6, for example, you get 873 hits. That's a bit better than it was in the old days. :-)
- When considering which of those books might work for you, recognize that some are written for specific audiences (like programmers), and those sort of books may not meet your particular needs (unless you're a coder and you're trying to come up to speed for IPv6).
- Also recognize that IPv6 is rapidly evolving, so beware of any books that haven't been recently updated.

Hurricane Electric's IPv6 Certification



Hurricane Electric Free IPv6 Certification

http://ipv6.he.net/certification/

Hurricane Electric Free IPv6 Certification

IPv6 Certifications

Welcome to the Hurricane Electric IPv6 Certification Project. This tool will allow you to certify your ability to configure IPv6, and to validate your IPv6 servers configuration.

Through this test set you will be able to:

- Prove that you have IPv6 connectivity
- Prove that you have a working IPv6 web server
- Prove that you have a working IPv6 email address
- Prove that you have working forward IPv6 DNS
- Prove that you have working reverse IPv6 DNS for your mail server
- Prove that you have name servers with IPv6 addresses that can respond to queries via IPv6
- Prove your knowledge of IPv6 technologies through quick and easy testing

You will also demonstrate that you are familiar with IPv6 concepts such as:

- the format of IPv6 addresses
- AAAA records
- reverse DNS for IPv6
- the IPv6 localhost address
- the IPv6 default route
- the IPv6 documentation prefix
- the IPv6 link local prefix
- the IPv6 multicast prefix
- how to do an IPv6 ping

Thanks For the Chance To Talk!

- Are there any questions?